Learning neural networks within Jupyter Notebook (part 1/2)

Piotr Migdał
https://p.migdal.pl

PiterPy
2-3 Nov 2018, Saint Petersburg

PhD in quantum physics theory
(2014, ICFO, Barcelona)

data science consultant

machine learning
deep learning
interactive data visualisations

$$i\hbar\dot{\psi} = \left(-\frac{\hbar^2}{2m}\nabla^2 + V(x)\right)\psi$$

$$\hat{H}\psi = E\psi$$

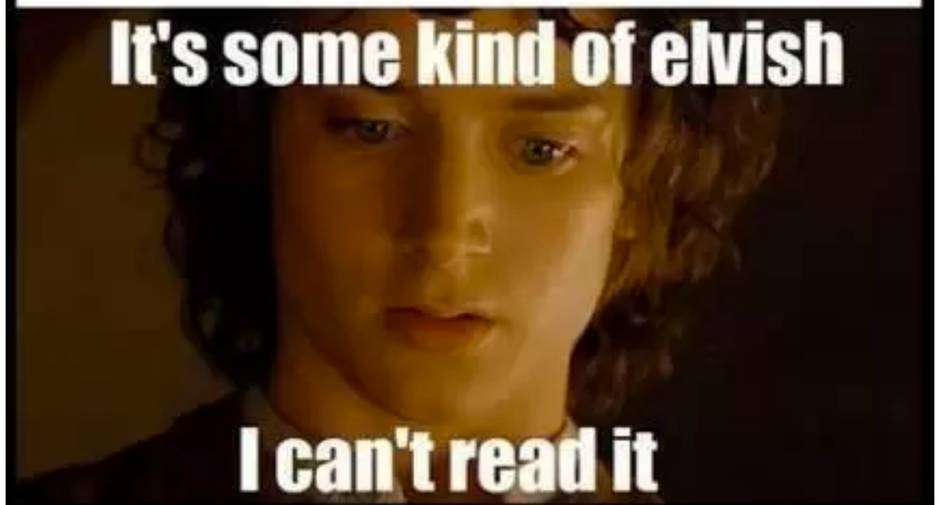

http://p.migdal.pl/2016/08/15/quantum-mechanics-for-high-school-students.html

Quantum Game with Photons

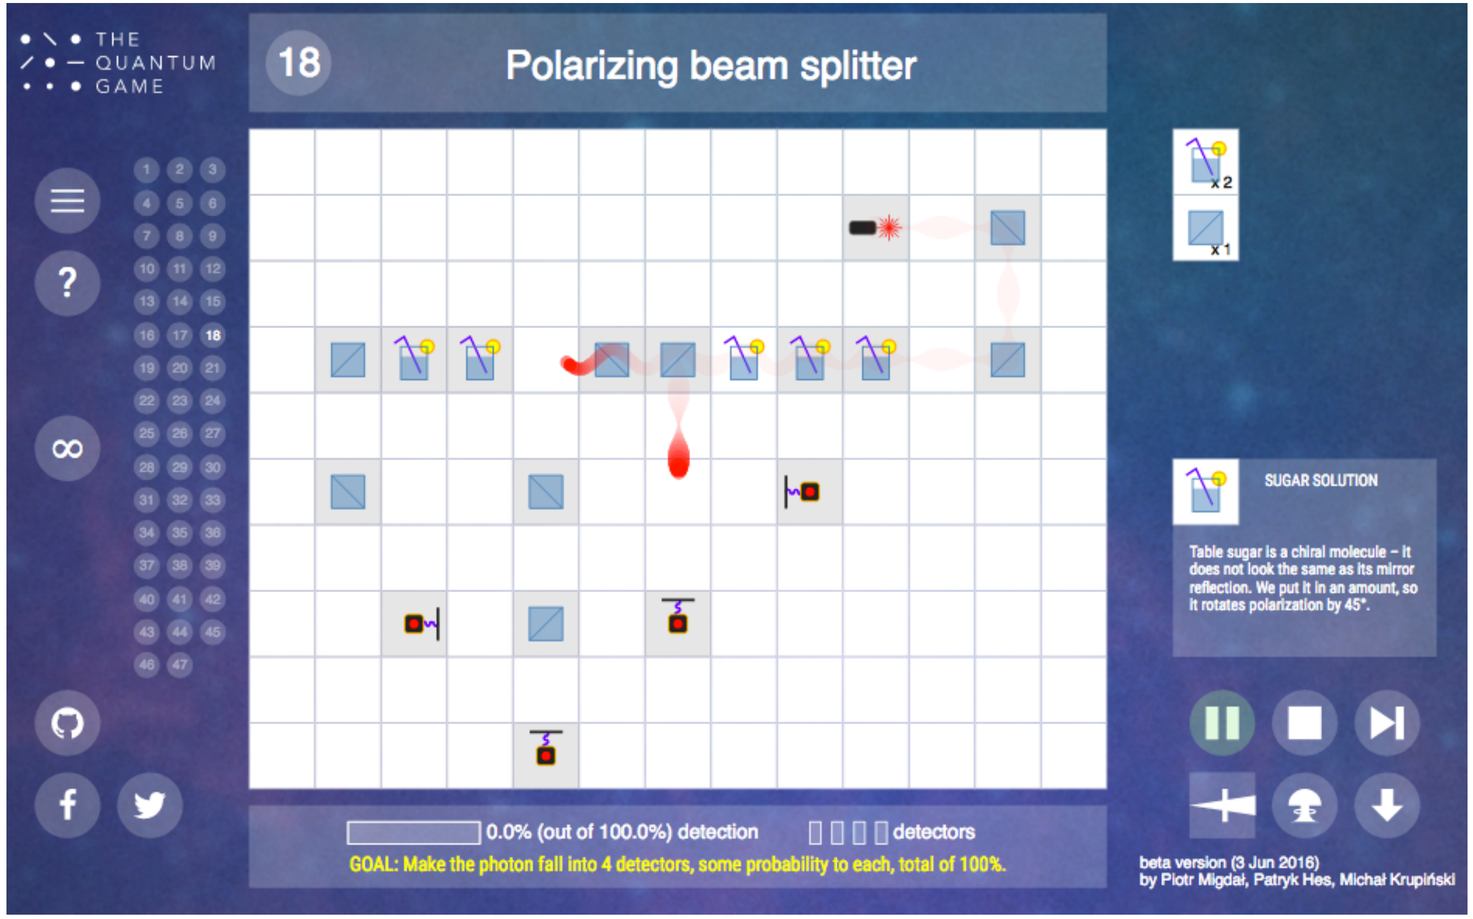

http://quantumgame.io

https://github.com/stared/science-based-games-list - more games

Deep learning
is powerful

AlphaGo

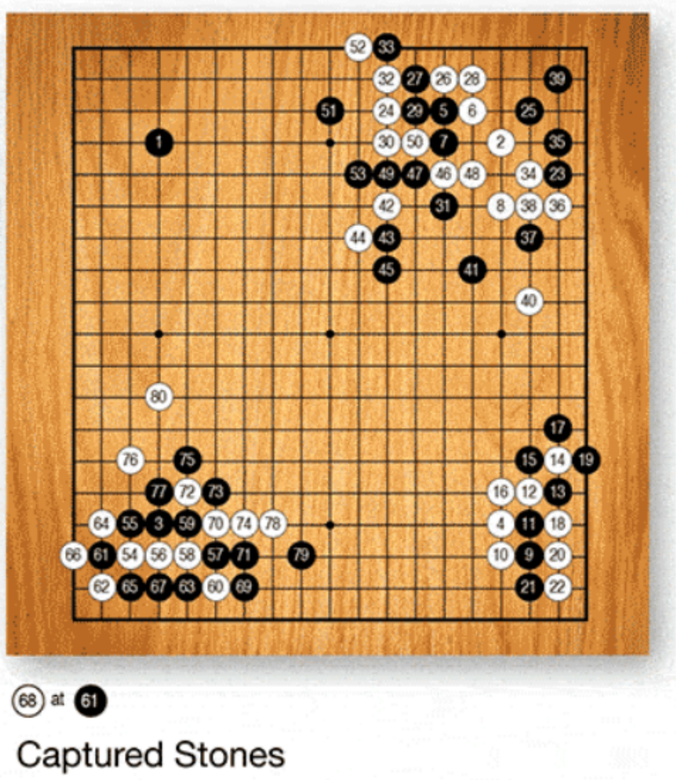

(68) at (61)

Captured Stones

https://deepmind.com/blog/alphago-zero-learning-scratch/

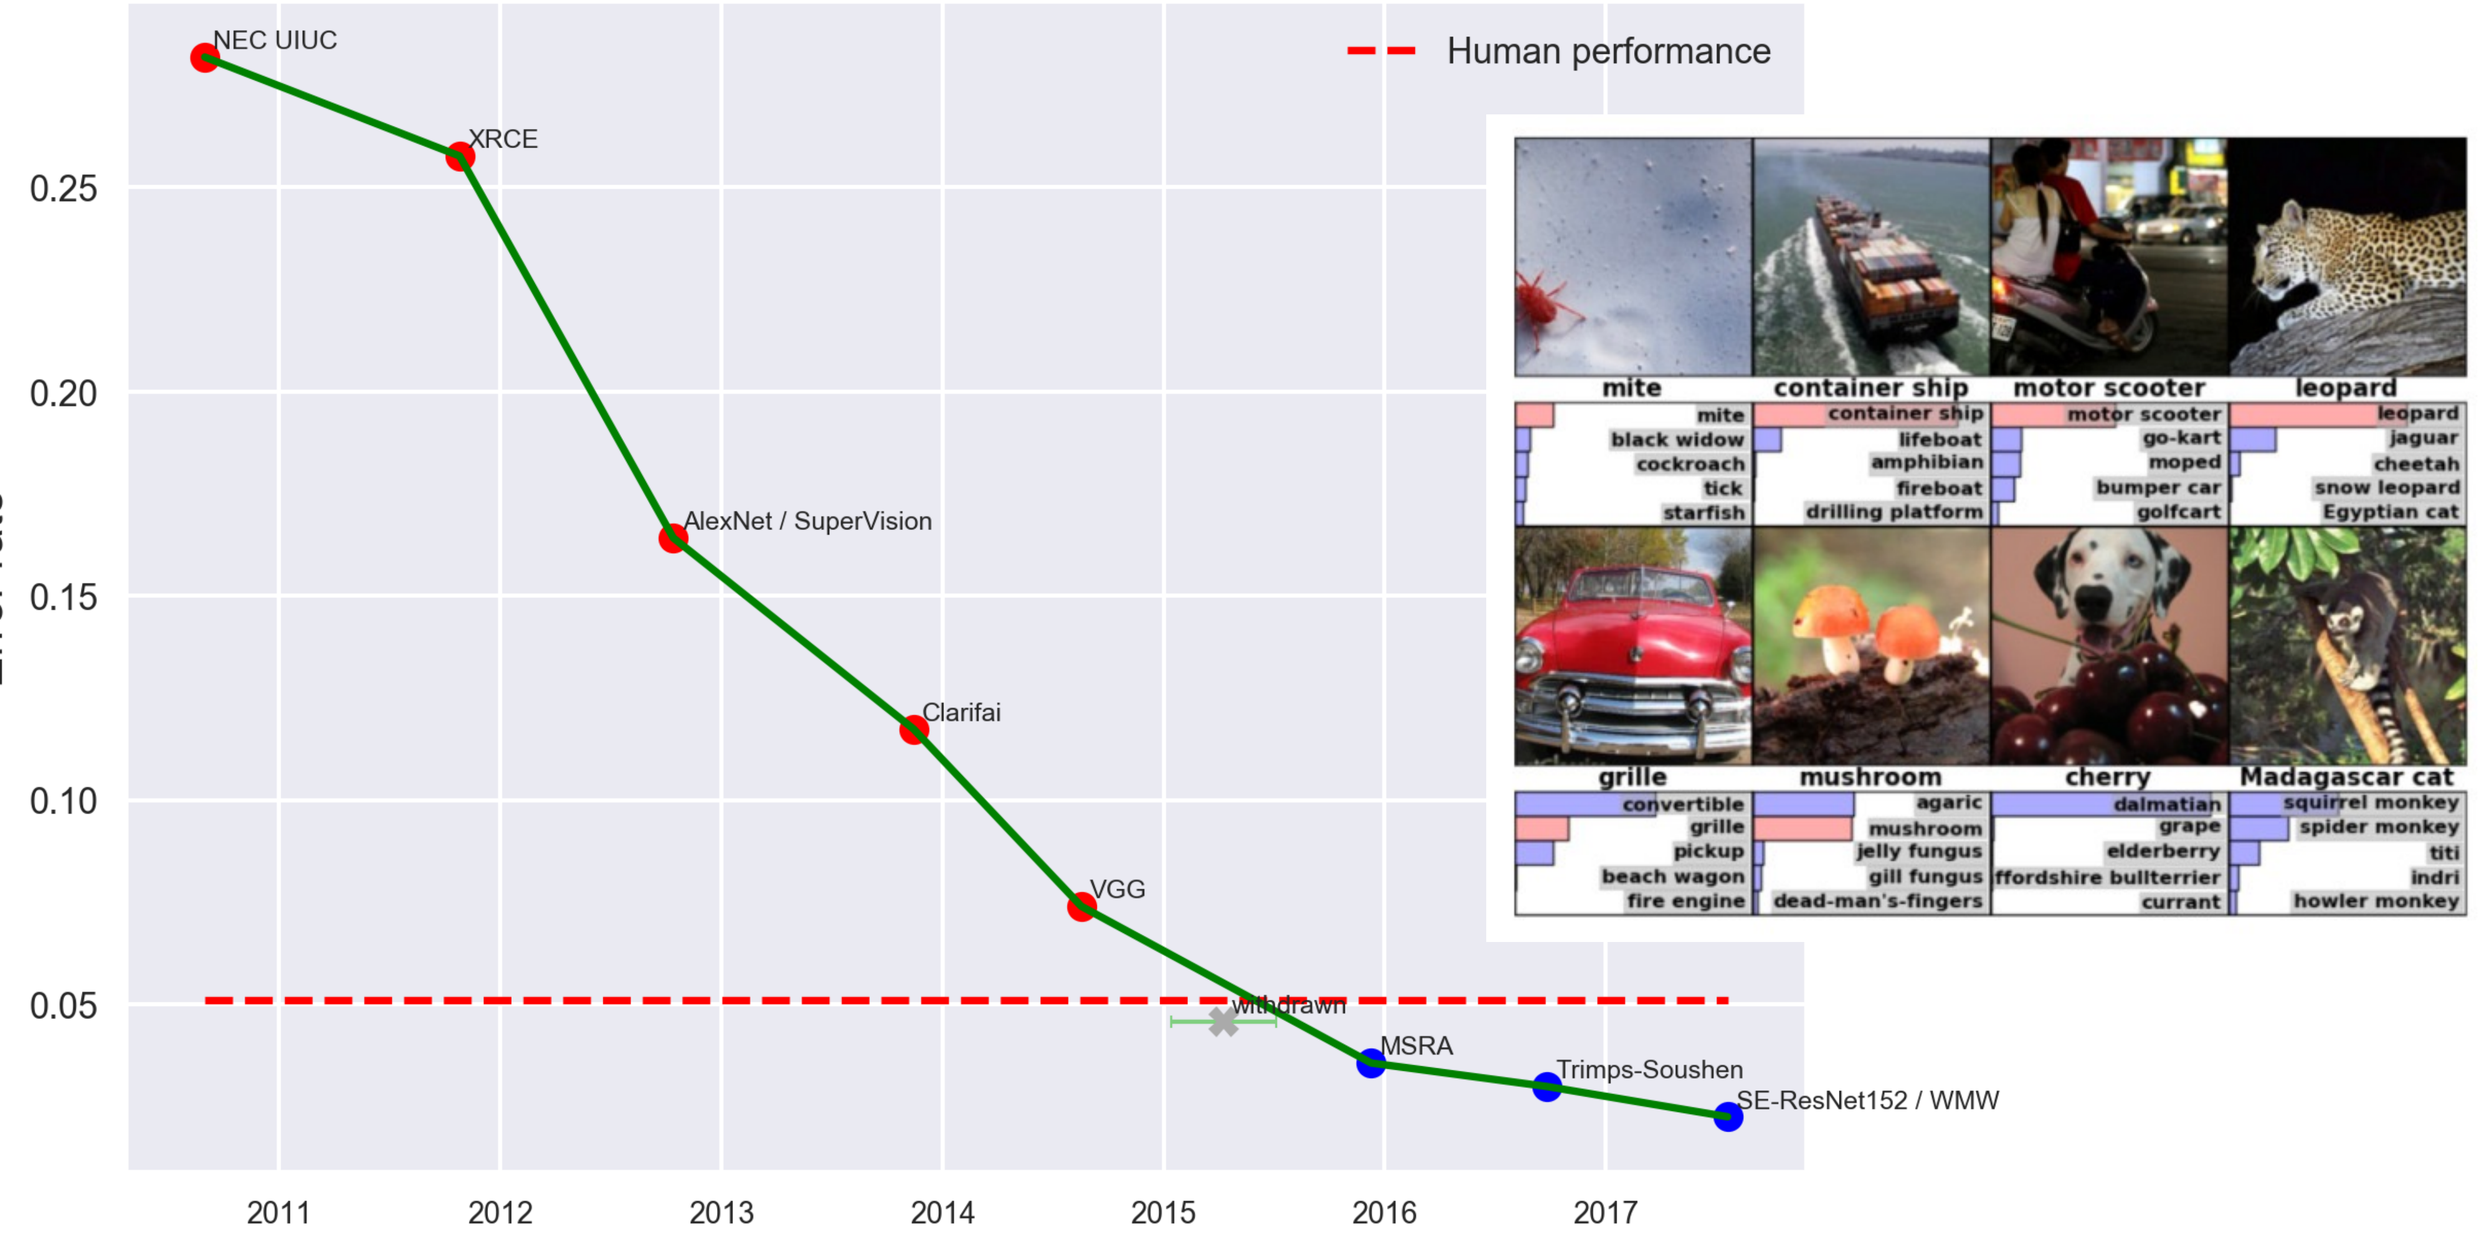

https://www.eff.org/ai/metrics

What do they see?

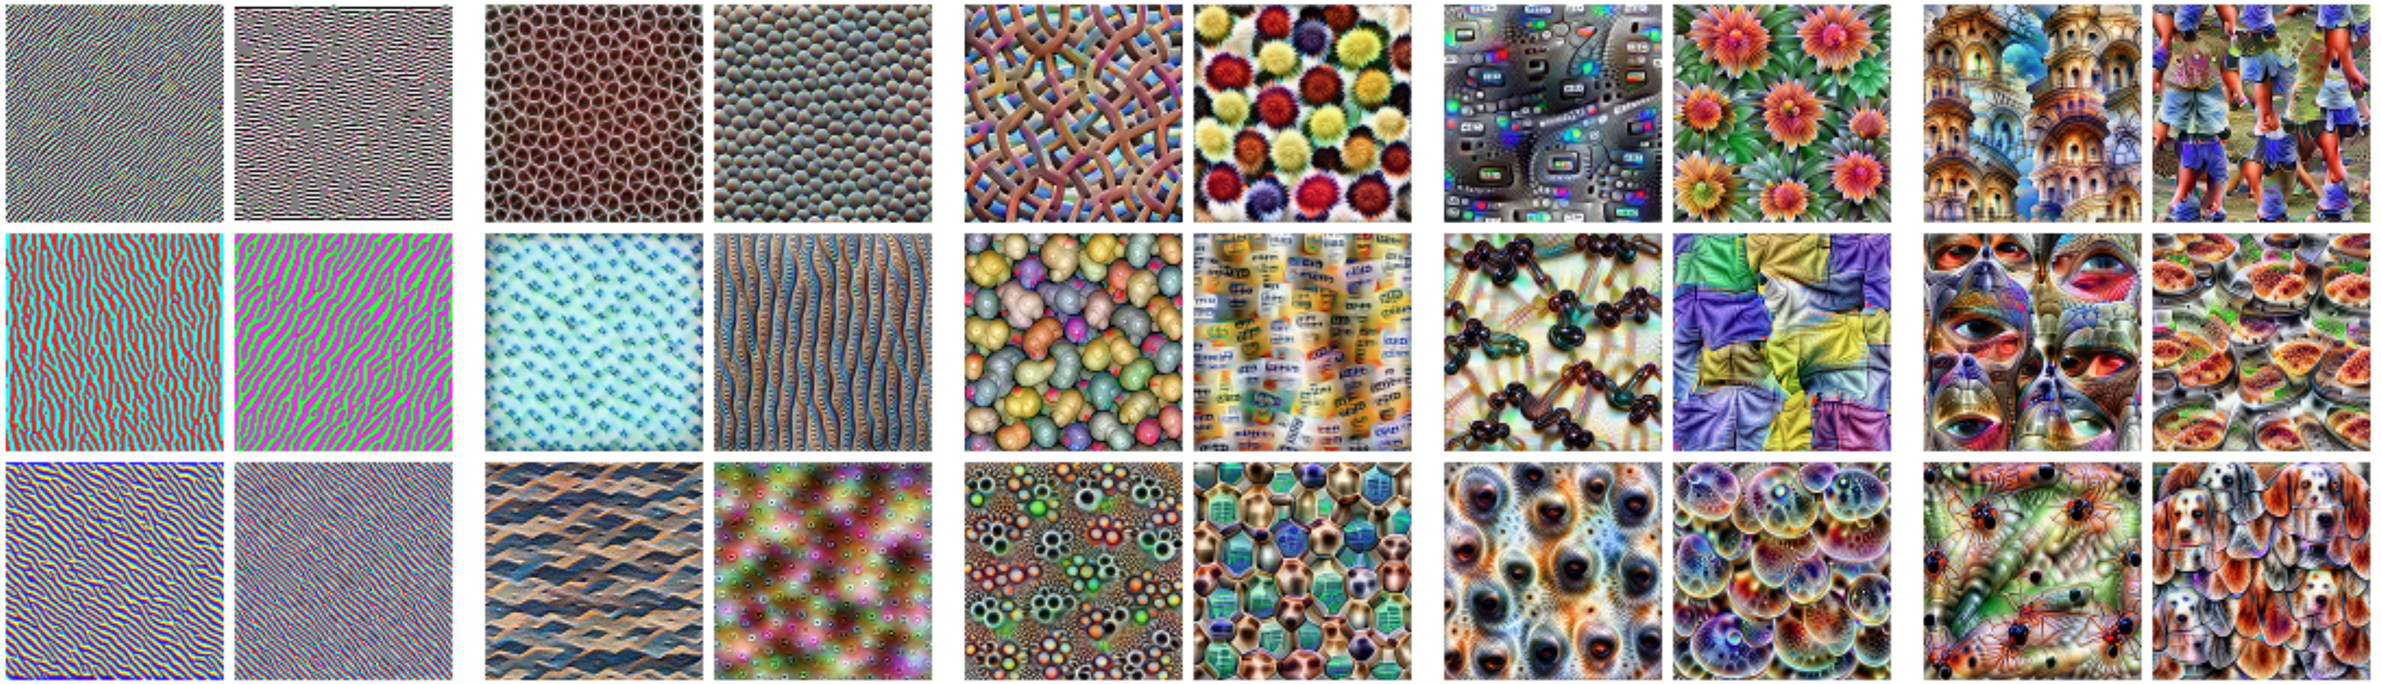

Edges (layer conv2d0) **Textures** (layer mixed3a) **Patterns** (layer mixed4a) **Parts** (layers mixed4b & mixed4c) **Objects** (layers mixed4d & mixed4e)

Feature visualization allows us to see how GoogLeNet [1], trained on the ImageNet [2] dataset, builds up its understanding of images over many layers. Visualizations of all channel are available in the appendix.

https://distill.pub/2017/feature-visualization/
https://tools.google.com/seedbank/seed/5086441721823232

What do they see?

Neuron

`layer`$_n$`[x,y,z]`

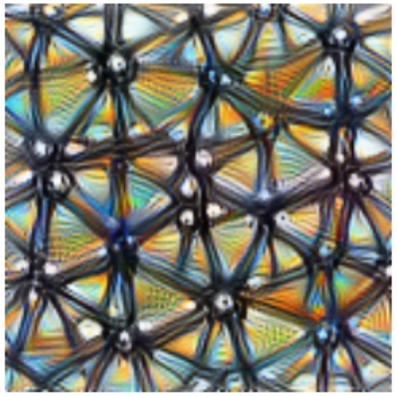

Channel

`layer`$_n$`[:,:,z]`

Layer/DeepDream

`layer`$_n$`[:,:,:]`^{2}

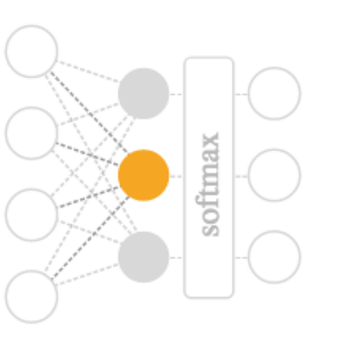

Class Logits

`pre_softmax[k]`

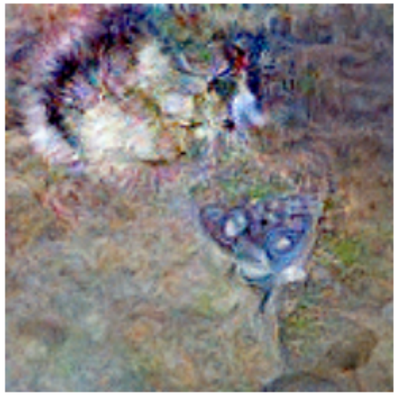

Class Probability

`softmax[k]`

https://distill.pub/2017/feature-visualization/

research paper on arXiv	Aug ‘15
code on GitHub	Oct ‘15
online tools (e.g. deepart.io, PRISMA)	Dec ‘15
movies	Apr ‘16

http://arxiv.org/abs/1508.06576

Skin cancer detection

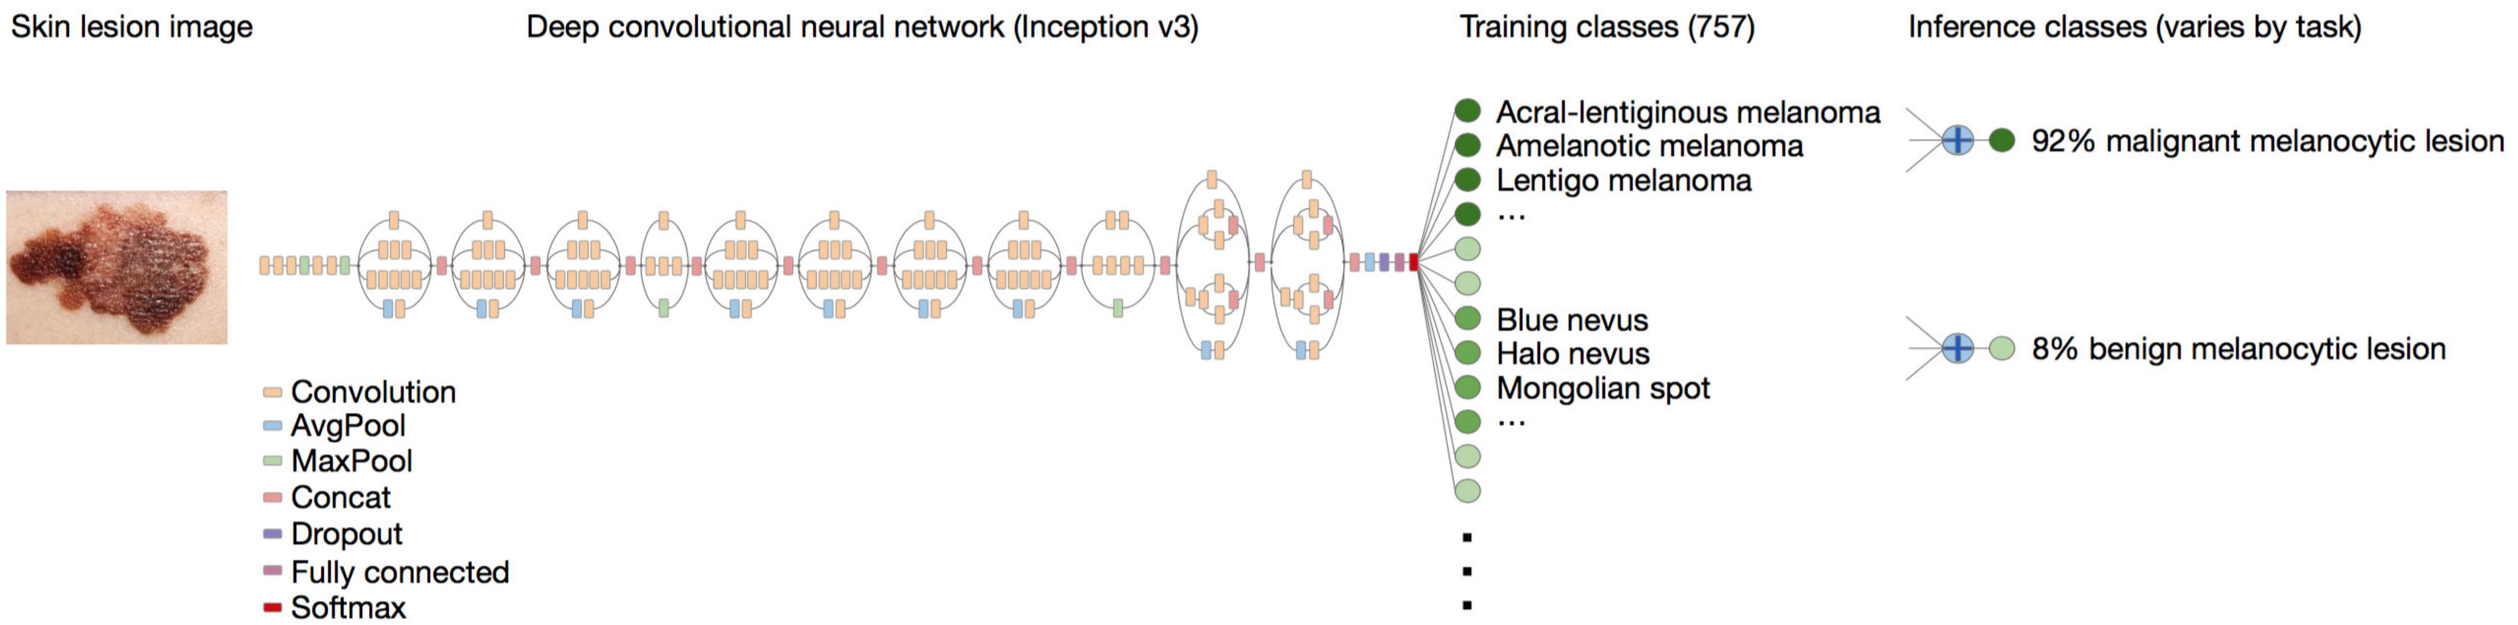

https://cs.stanford.edu/people/esteva/nature/

X-ray for detecting broken bones

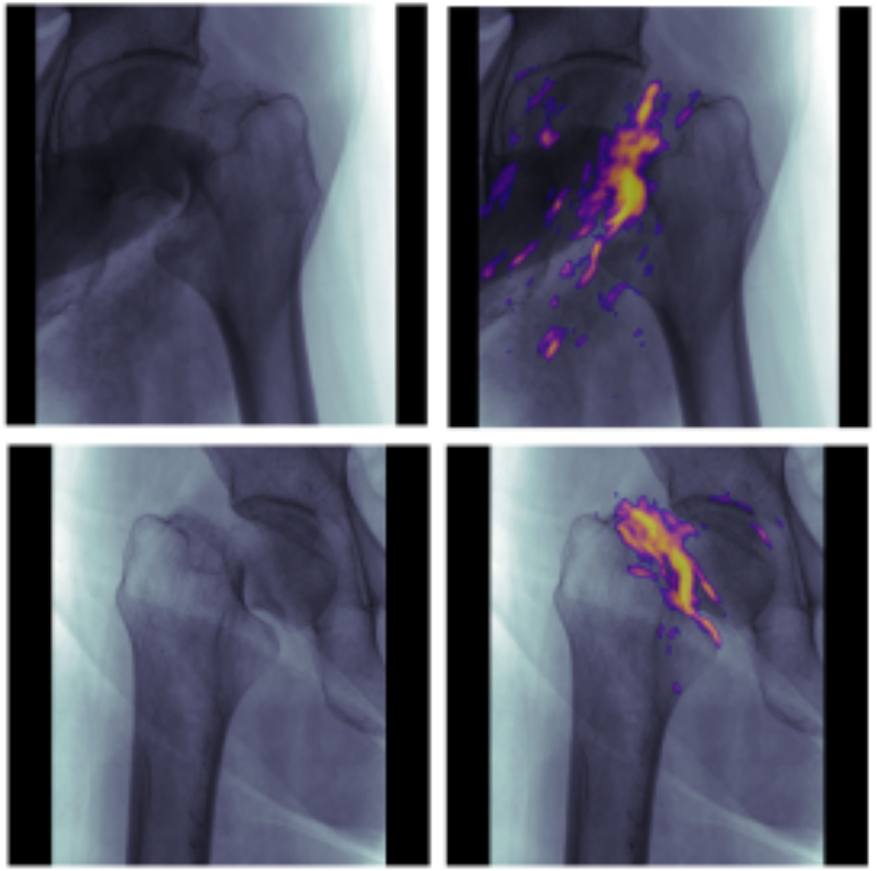

Original report
There is a transcervical fracture through the left neck of femur.

Generated report
There is a moderately displaced subcapital fracture of the left neck of femur.

Original report
Subcapital fracture of the right neck of femur with associated shortening.

Generated report
There is a significantly displaced subcapital fracture of the right neck of femur.

https://lukeoakdenrayner.wordpress.com/2018/06/05/explain-yourself-machine-producing-simple-text-descriptions-for-ai-interpretability/

Satellite images to maps

https://deepsense.ai/deep-learning-for-satellite-imagery-via-image-segmentation/

Image segmentation

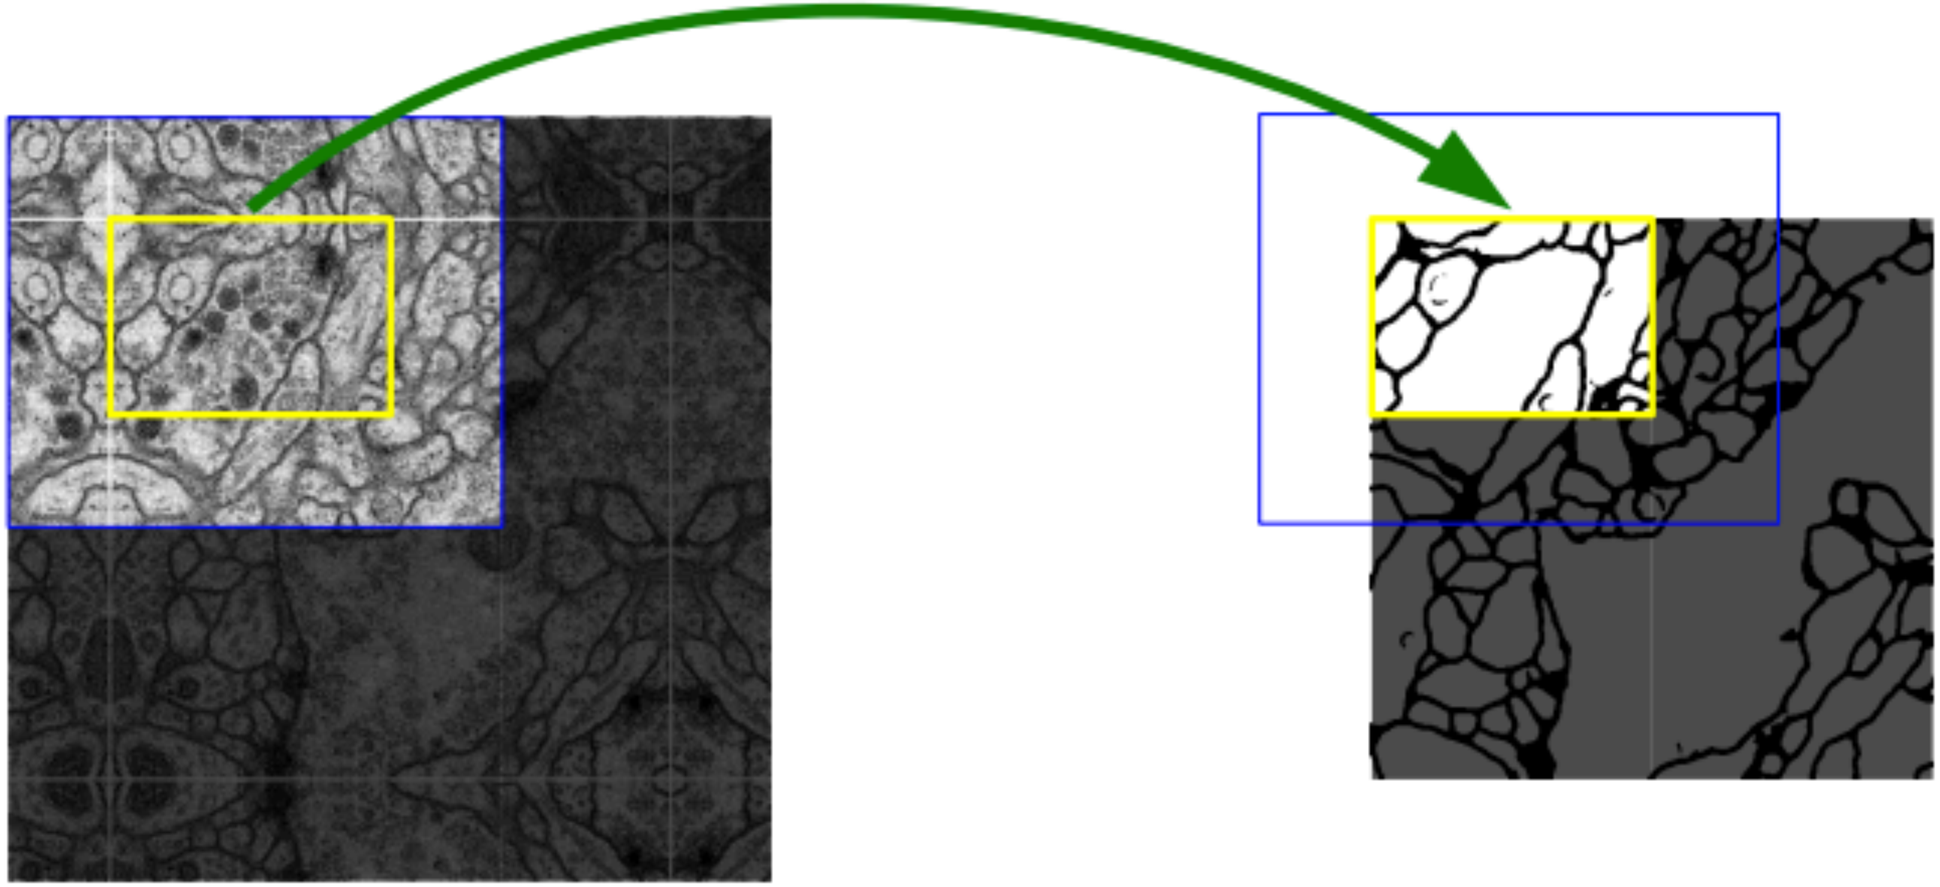

https://lmb.informatik.uni-freiburg.de/people/ronneber/u-net/

functional magnetic resonance -> picture in your mind

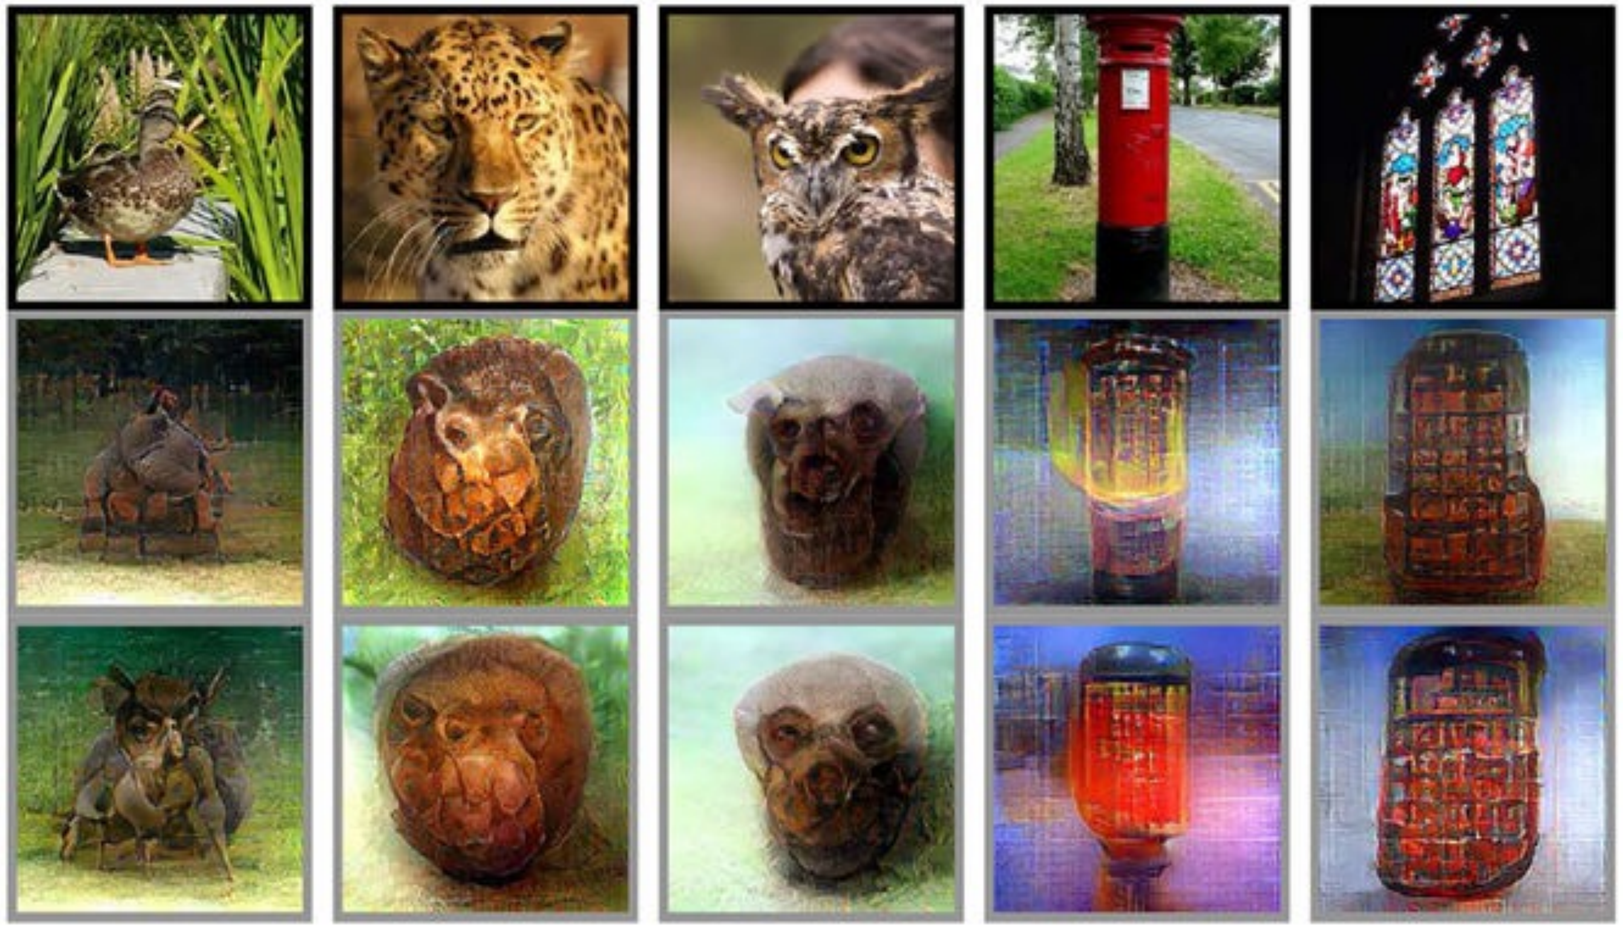

http://www.sciencemag.org/news/2018/01/mind-reading-algorithm-can-decode-pictures-your-head

So it is difficult, right?

Deep Learning

What society thinks I do

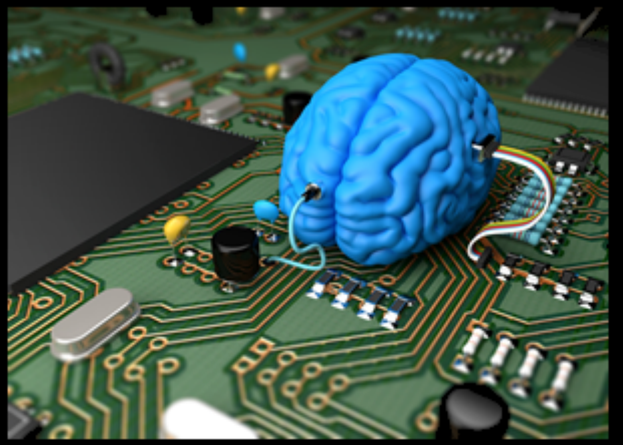

What my friends think I do

What other computer scientists think I do

What mathematicians think I do

What I think I do

```
In [1]:
import keras
Using TensorFlow backend.
```

What I actually do

Just learn it (pick Keras or PyTorch)

```
# create neural network architecture
model = Sequential()

model.add(Conv2D(32, (3, 3), activation='relu',
                 input_shape=(32, 32, 3)))
model.add(MaxPool2D())

model.add(Conv2D(64, (3, 3), activation='relu'))
model.add(MaxPool2D())

model.add(Flatten())
model.add(Dense(10))
model.add(Activation('softmax'))

model.compile(optimizer='adam',
              loss='categorical_crossentropy',
              metrics=['accuracy'])

model_summary(model)

# loading data
(x_train, y_train), (x_test, y_test) = load_cifar10()
```

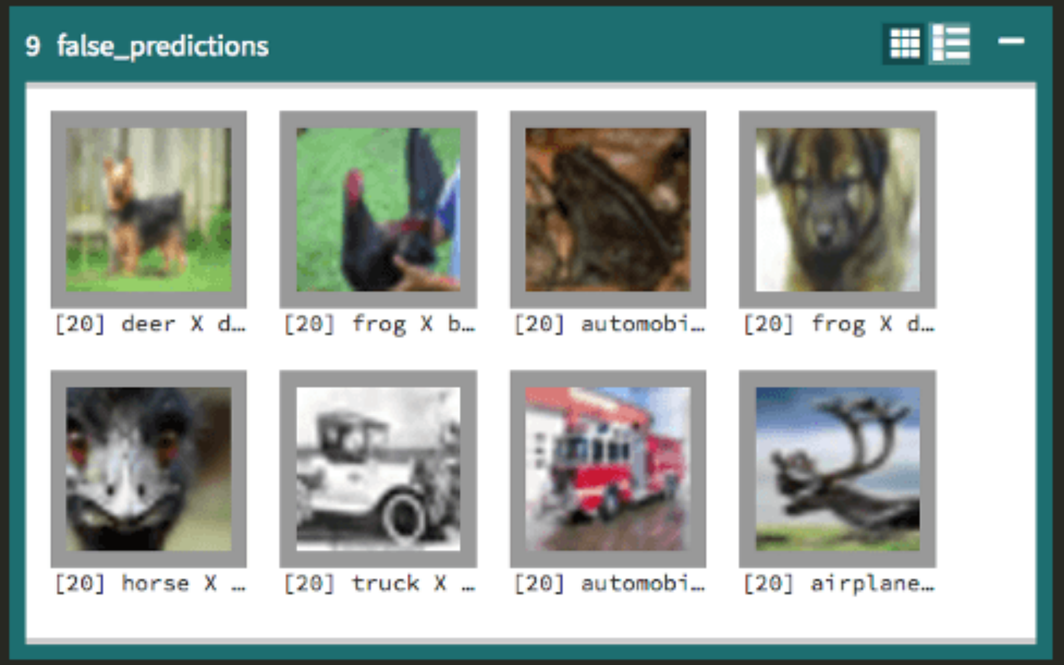

https://deepsense.ai/deep-learning-hands-on-image-classification/
https://deepsense.ai/keras-or-pytorch/

Other neural networks are fine

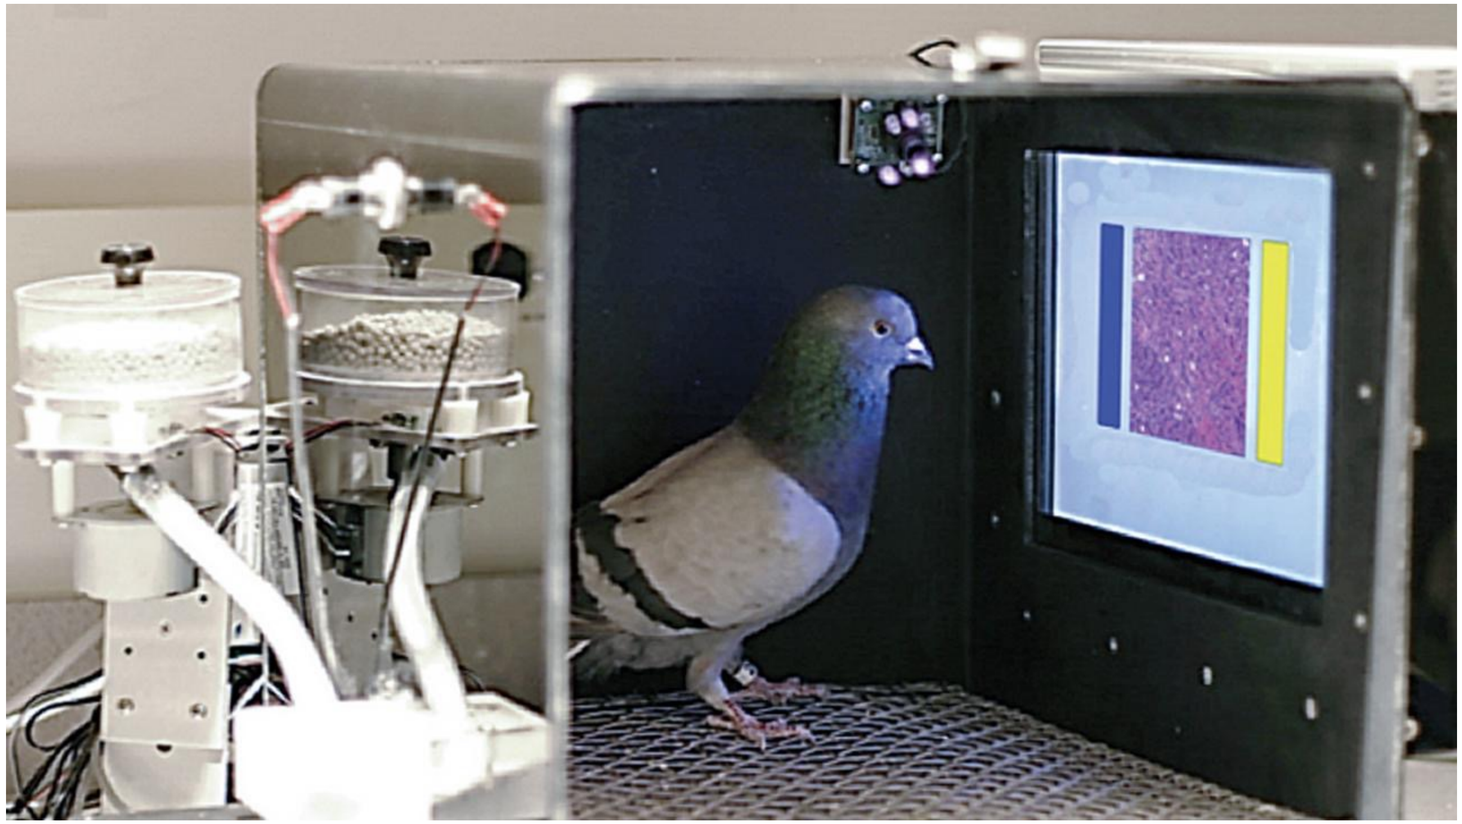

http://www.sciencemag.org/news/2015/11/pigeons-spot-cancer-well-human-experts

https://www.youtube.com/watch?v=fIzGjnJLyS0

And interactively

Decision trees, visually

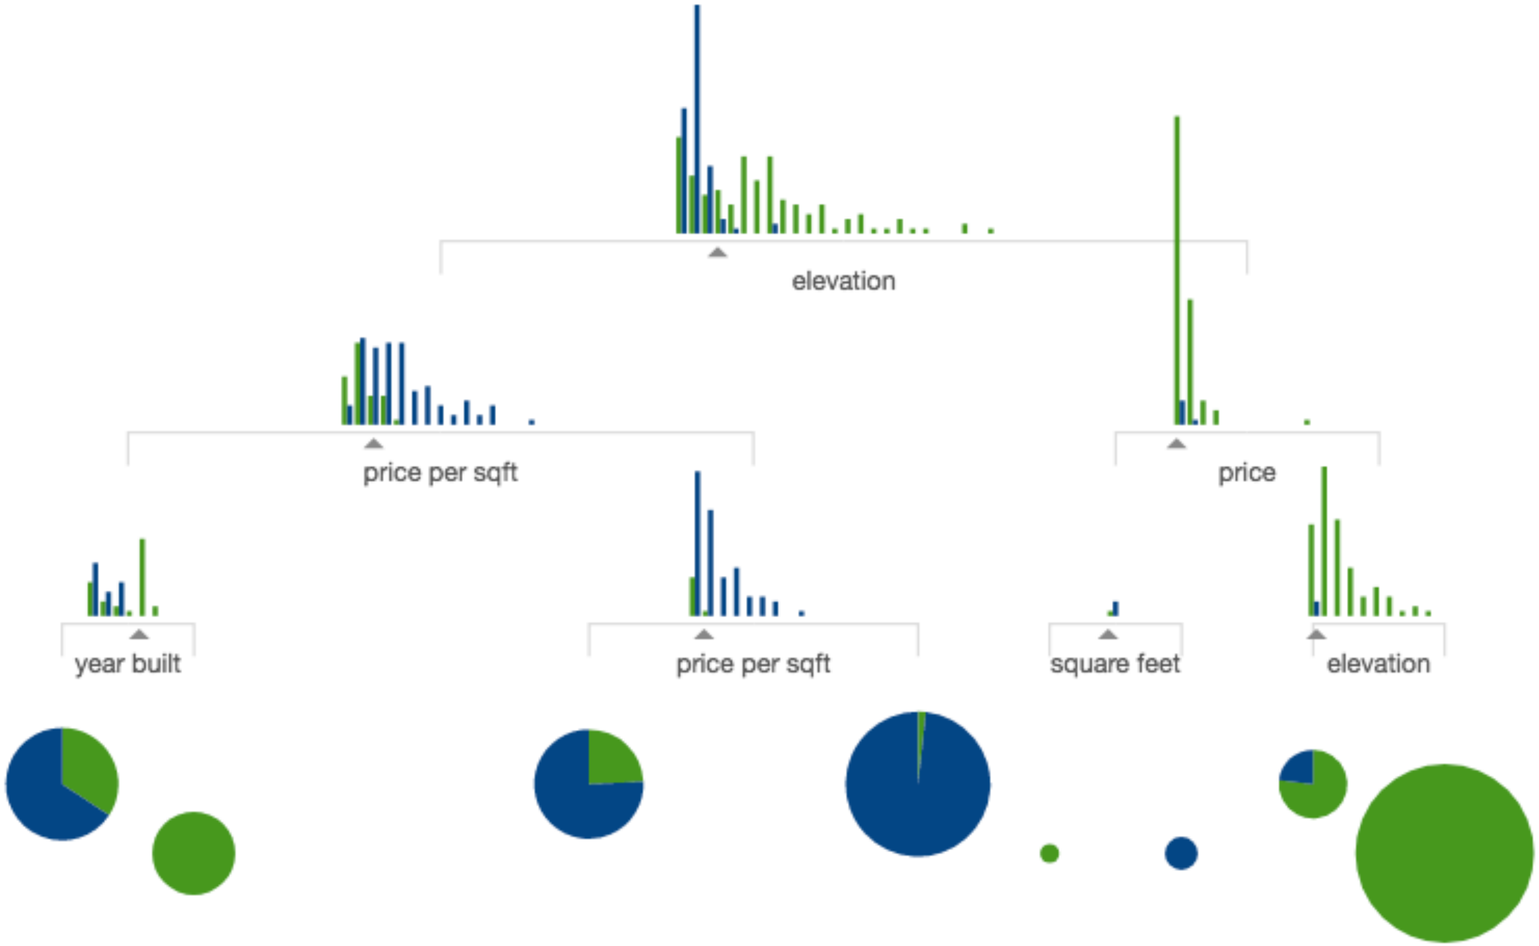

http://www.r2d3.us/visual-intro-to-machine-learning-part-1/

Interactive ML

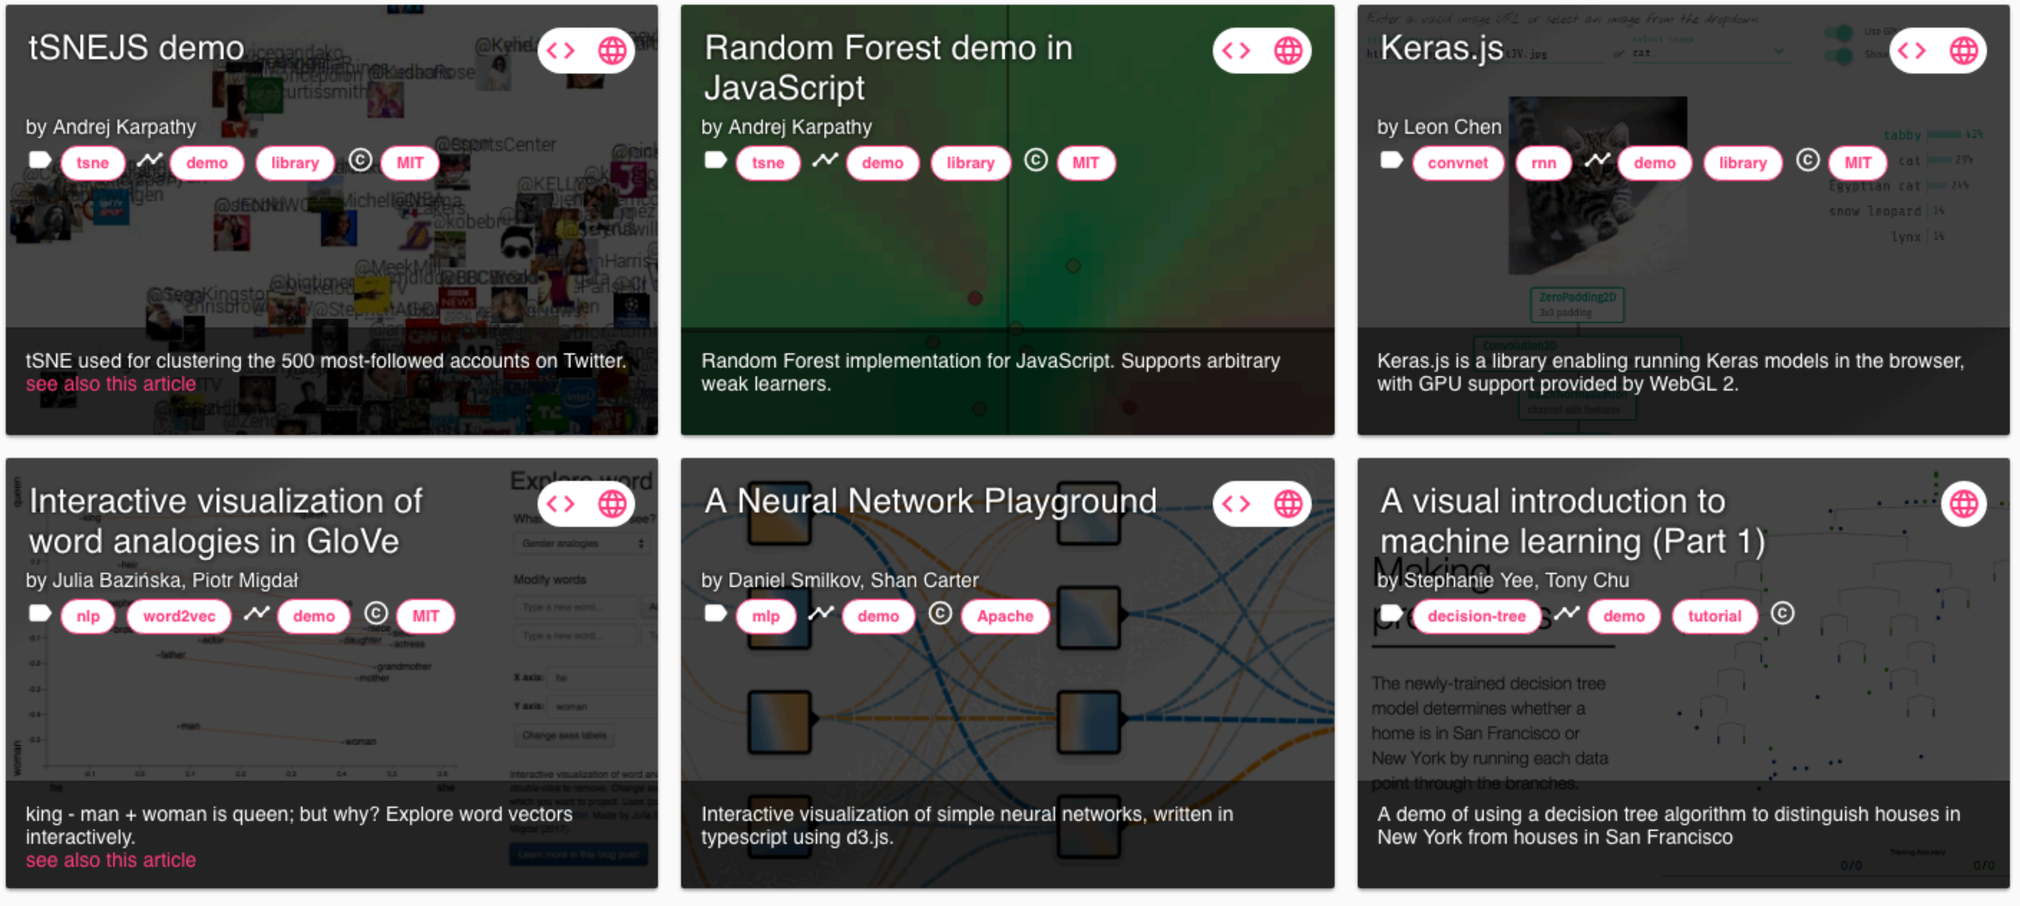

https://p.migdal.pl/interactive-machine-learning-list/

https://transcranial.github.io/keras-js/#/mnist-cnn

Generating faces

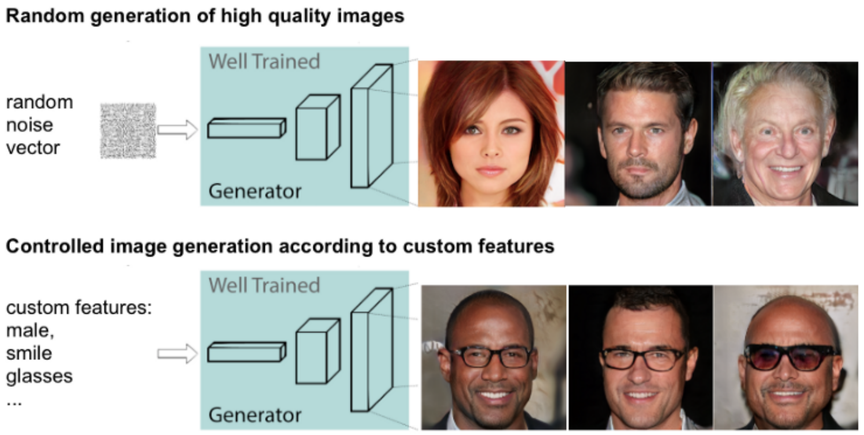

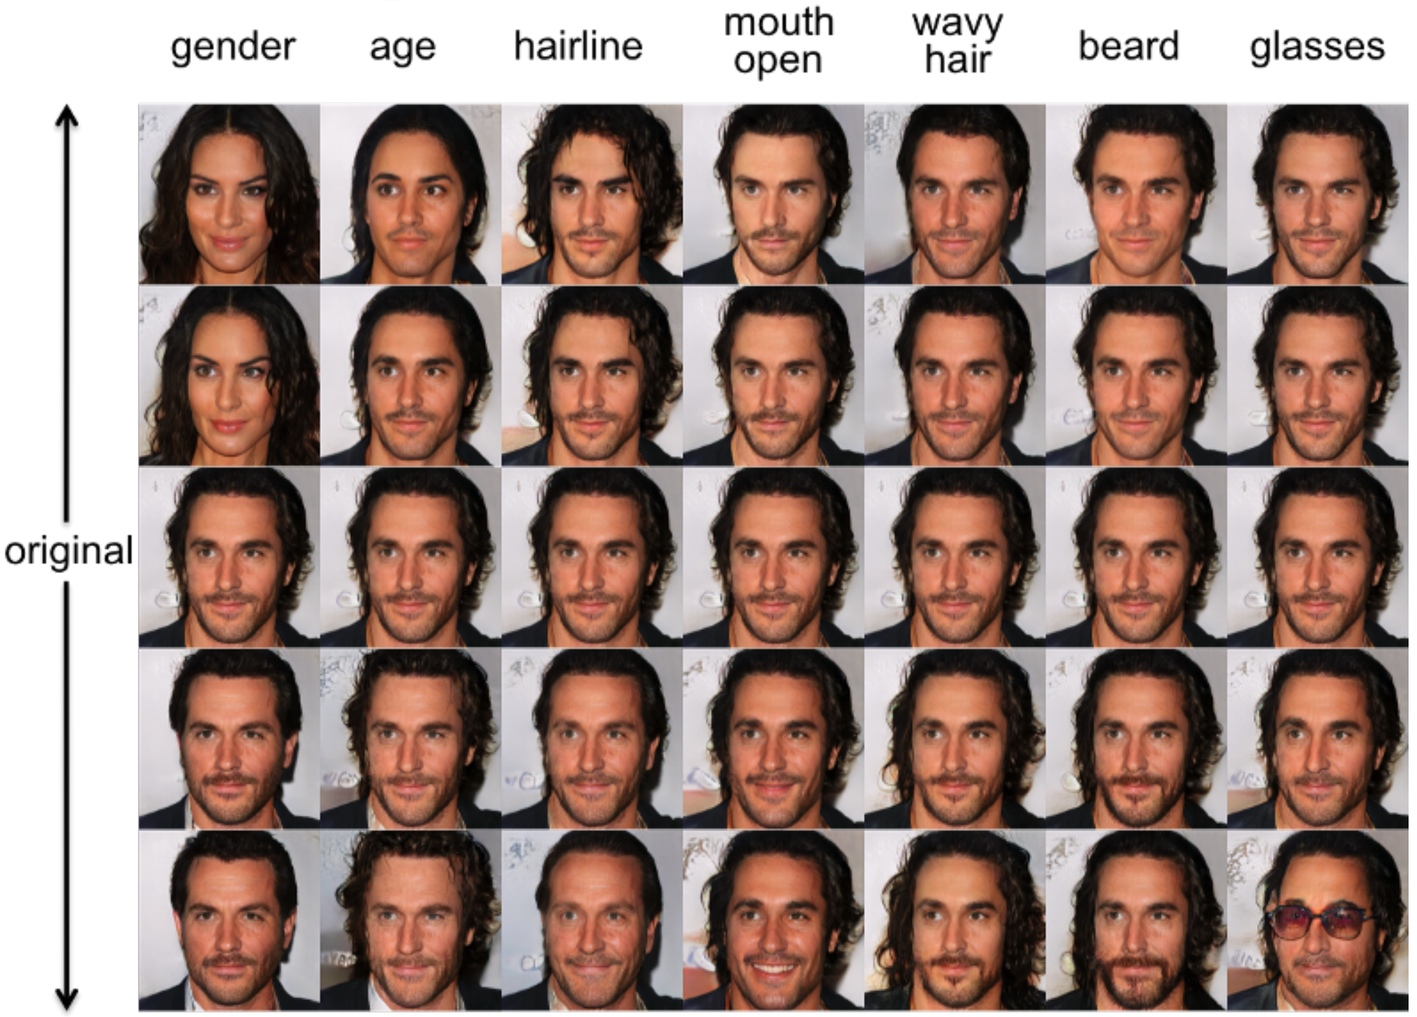

https://blog.insightdatascience.com/generating-custom-photo-realistic-faces-using-ai-d170b1b59255

Generating faces… online

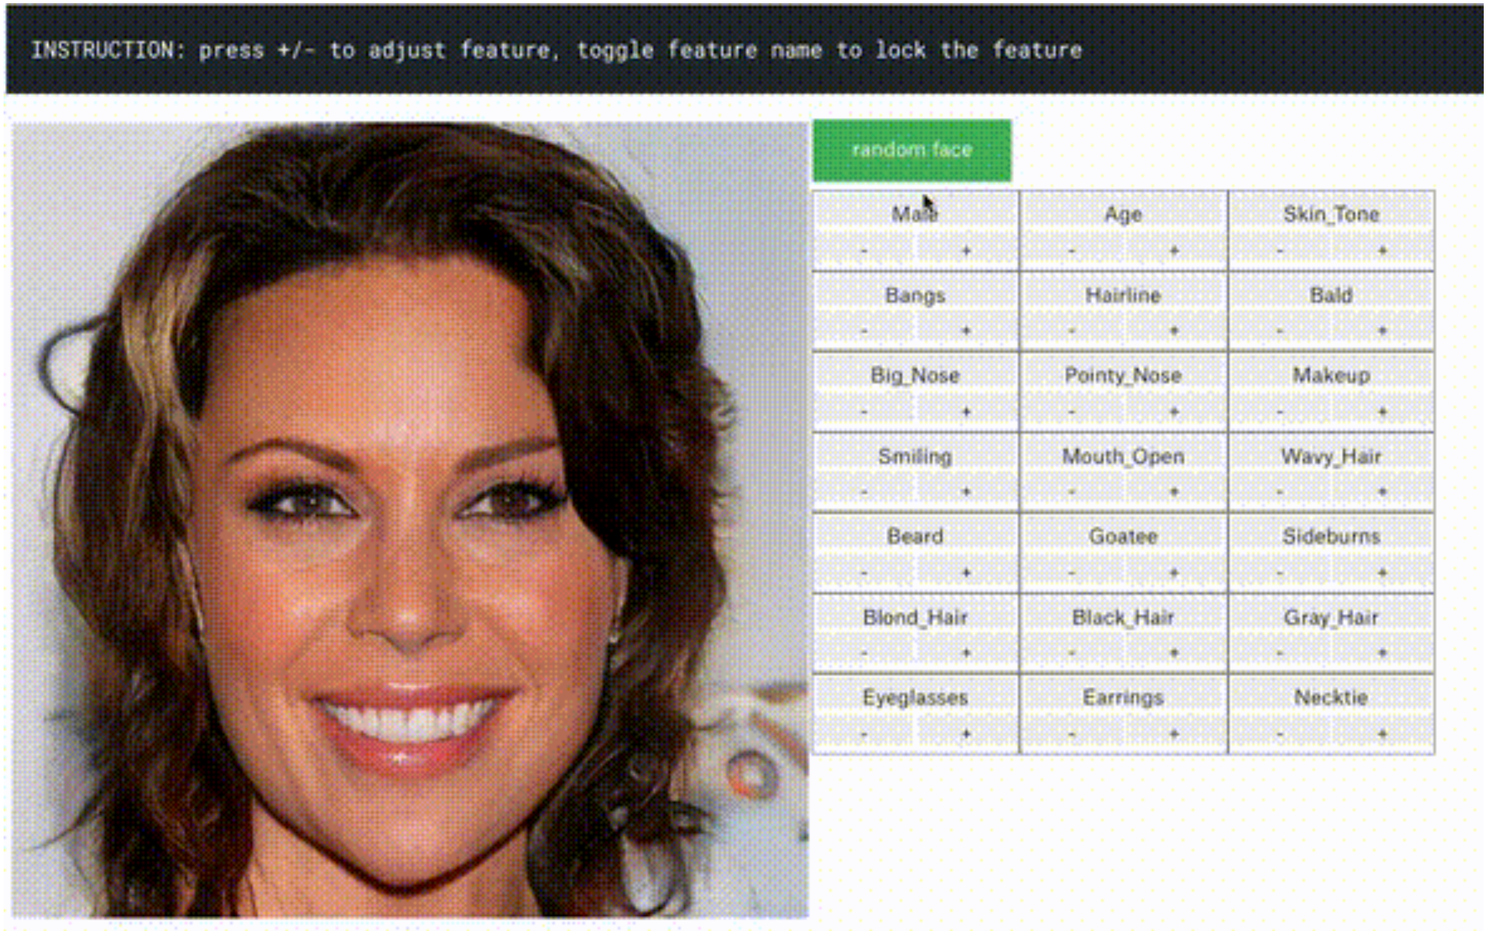

https://blog.insightdatascience.com/generating-custom-photo-realistic-faces-using-ai-d170b1b59255
play interactively in a Jupyter Notebook: https://www.kaggle.com/summitkwan/tl-gan-demo

Let’s learn it!

Real data, plots > arrays

```
In [10]: # feature (variable) importance
         rfr.feature_importances_

Out[10]: array([ 0.34966593,  0.30402689,  0.00540573,  0.03778628,  0.01359446,
                 0.06506118,  0.11507385,  0.10938569])

In [11]: # but we want to plot it
         pd.Series(rfr.feature_importances_, index=X.columns).plot(kind="barh")

Out[11]: <matplotlib.axes._subplots.AxesSubplot at 0x111427290>
```

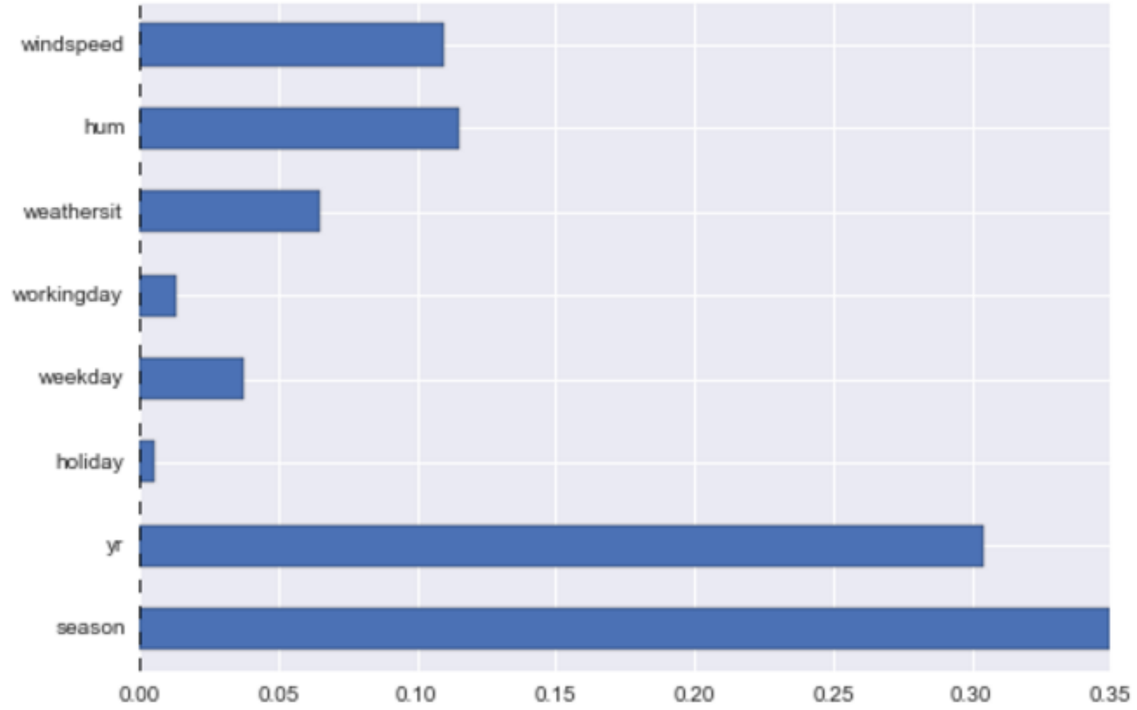

MNIST

MNIST

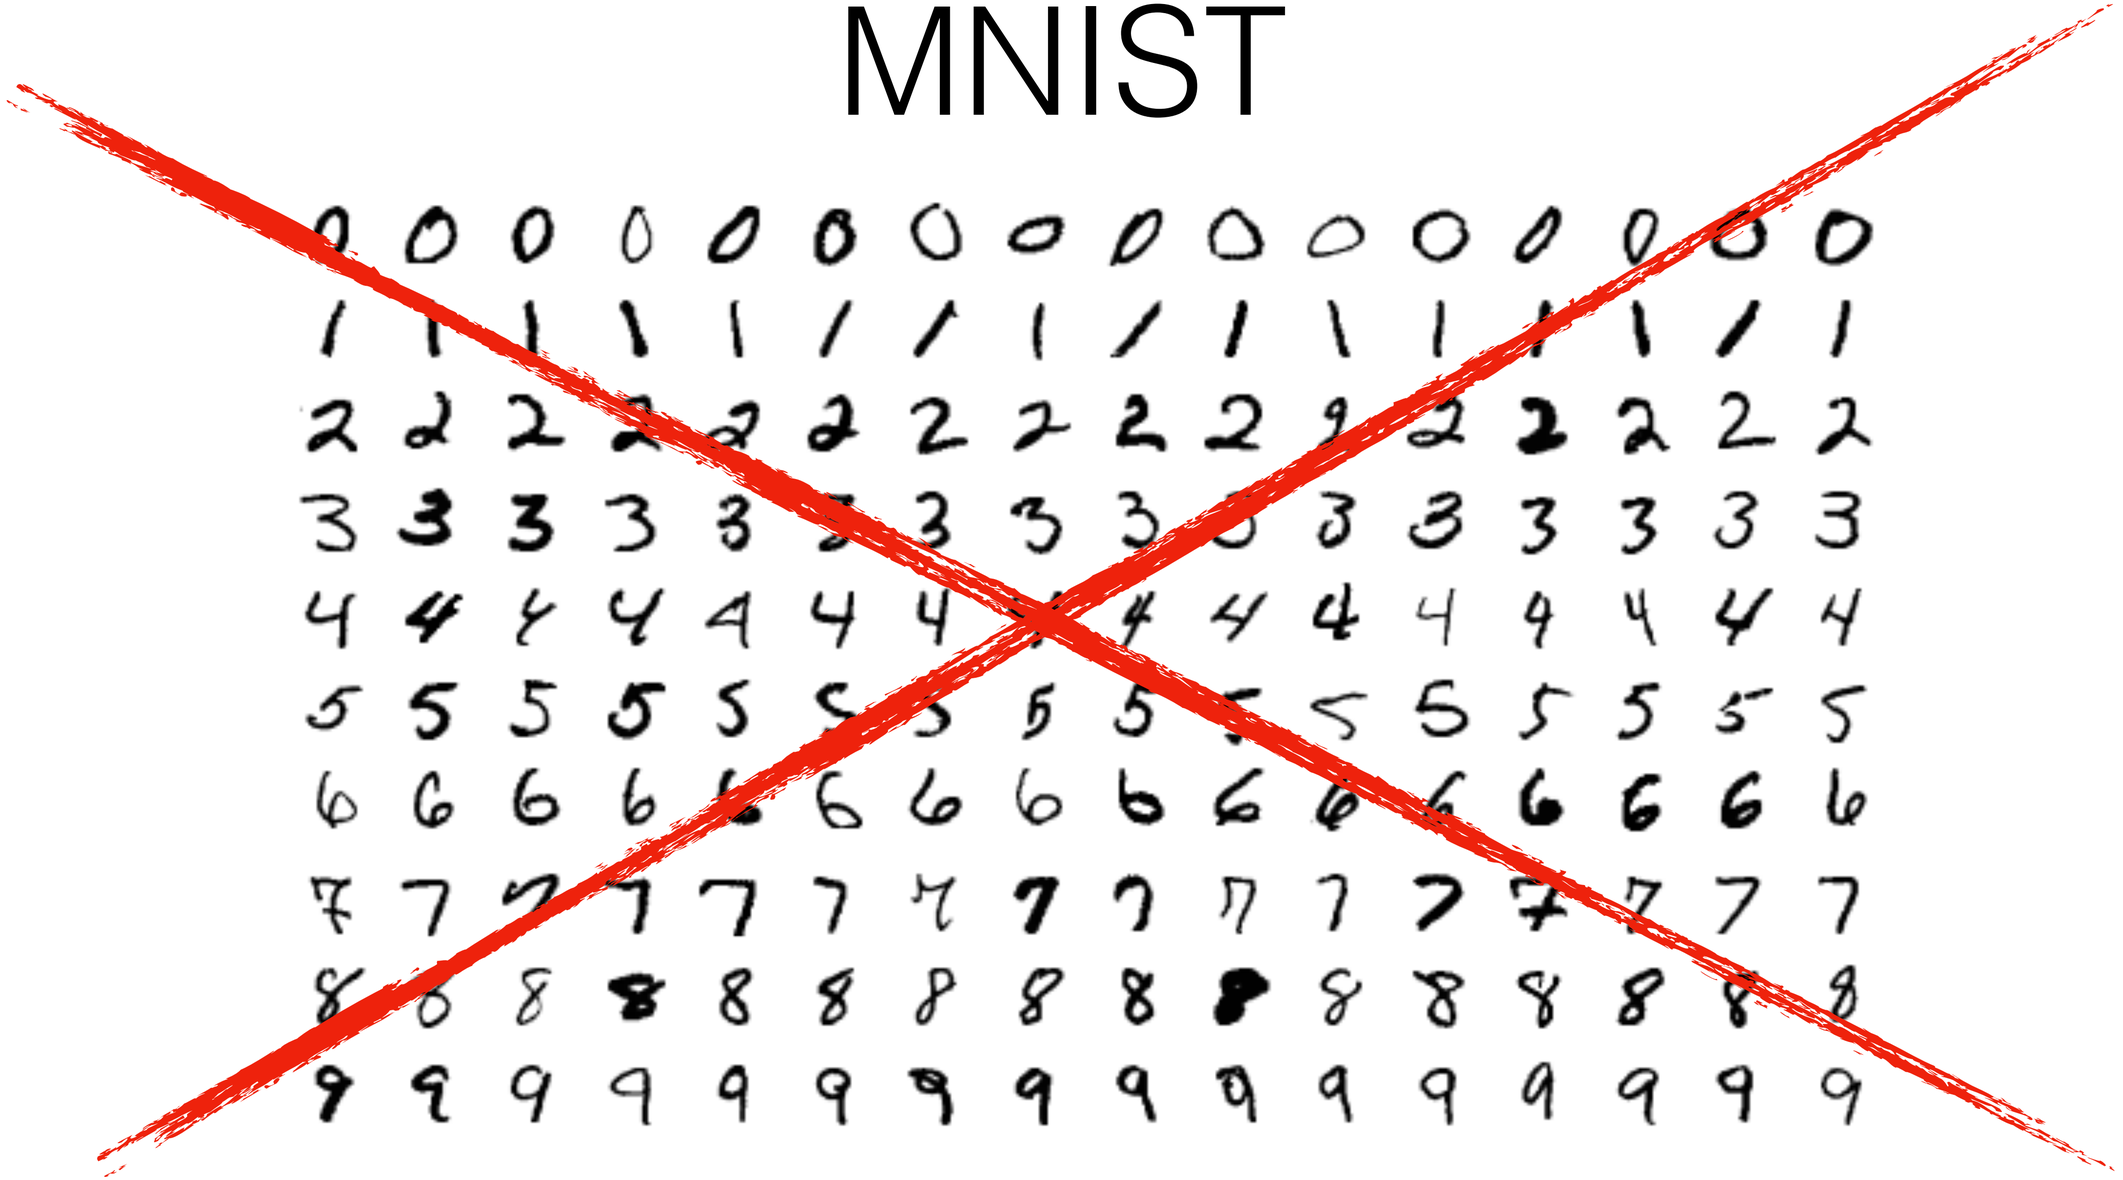

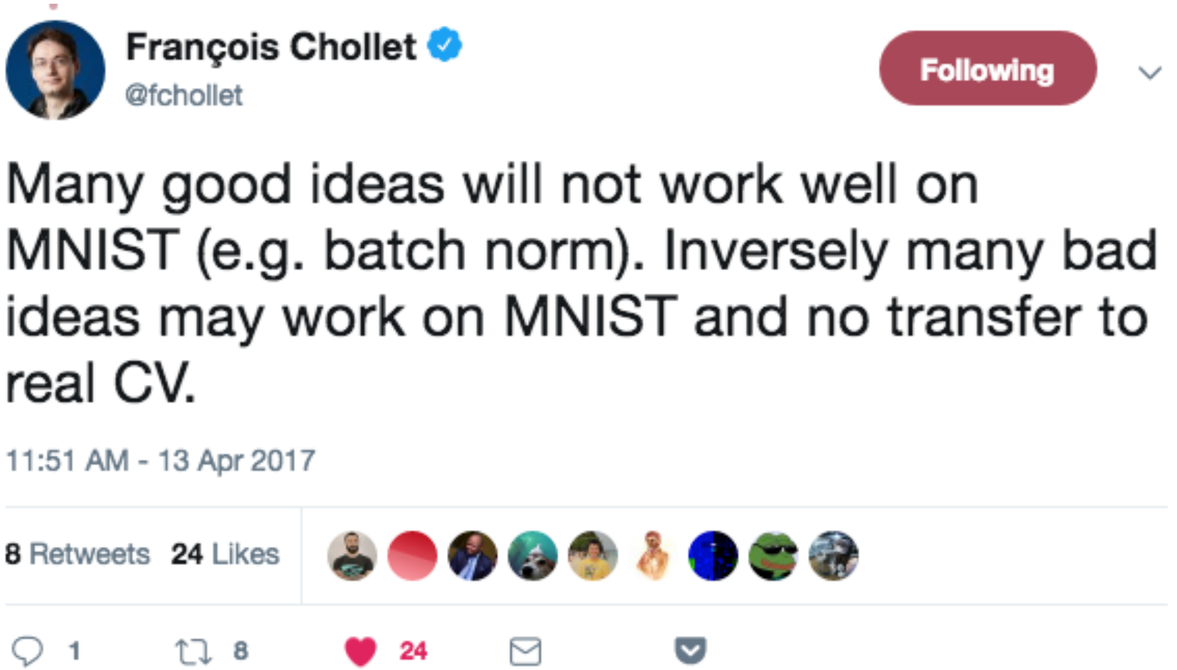

not-MNIST

fashion-MNIST

https://twitter.com/fchollet/status/852594987527045120
https://github.com/zalandoresearch/fashion-mnist

Ok, so maybe on a nicer dataset?

Drawing 5/6

Draw

angel

in under 20 seconds

Got It!

https://twitter.com/kcimc/status/902229612666658816 - more example of averages
https://medium.com/@enjalot/machine-learning-for-visualization-927a9dff1cab - Ian Johnson

https://github.com/stared/keras-interactively-piterpy2018

Thank you!
Questions?

blog:	p.migdal.pl
email:	pmigdal@gmail.com
twitter:	@pmigdal
In Browser AI:	inbrowser.ai / @InBrowserAI

Decision trees, visually

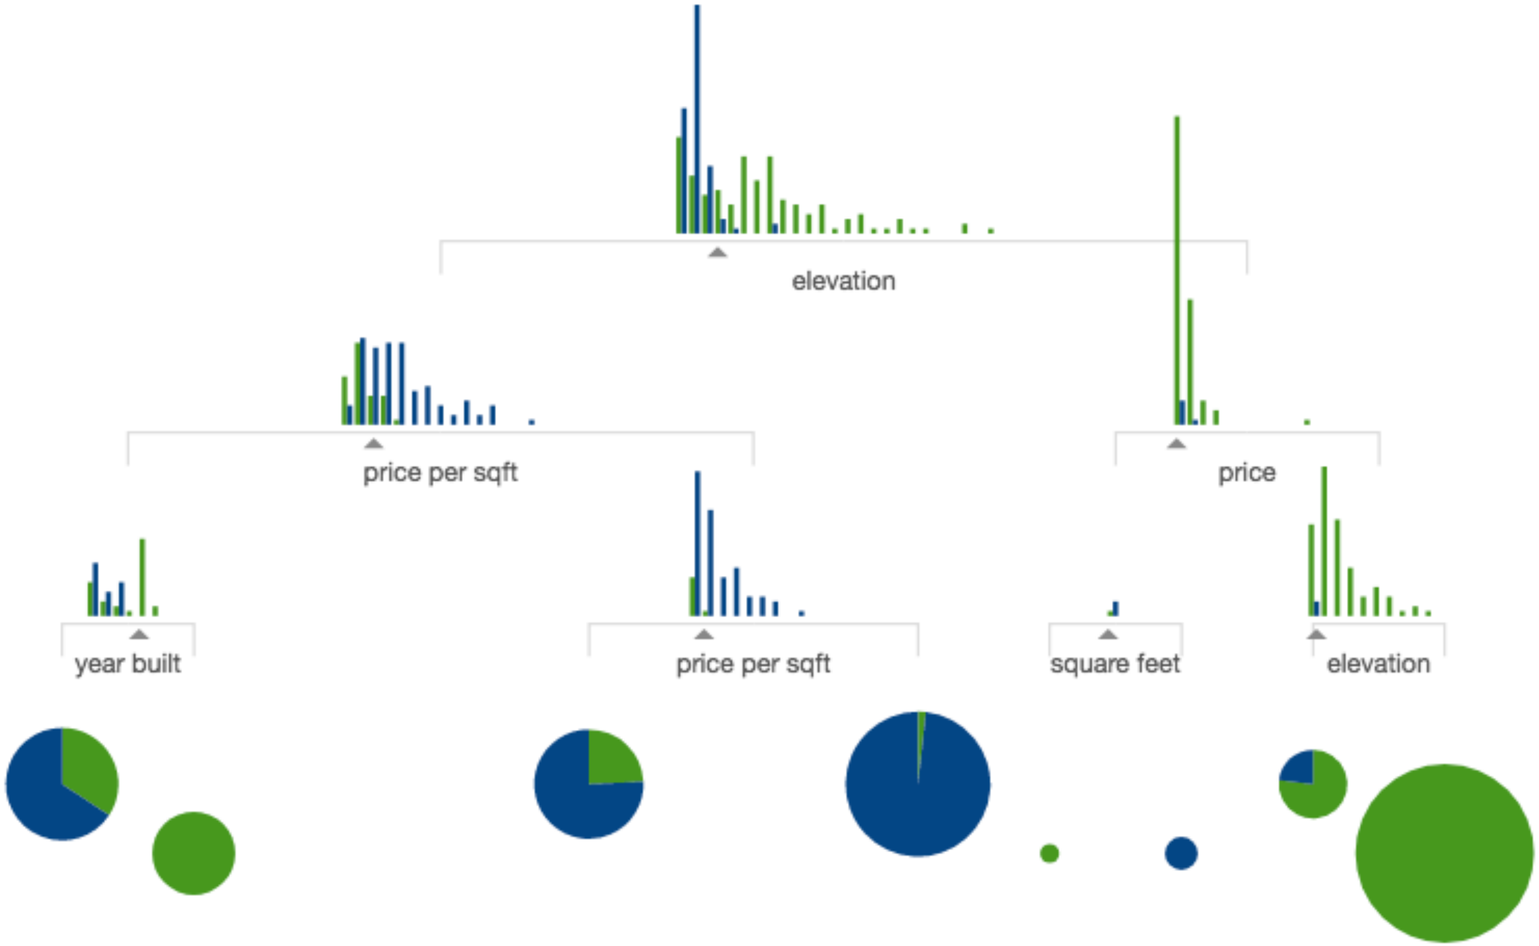

http://www.r2d3.us/visual-intro-to-machine-learning-part-1/

Thank you!
Questions?

blog:	p.migdal.pl
email:	pmigdal@gmail.com
twitter:	@pmigdal
In Browser AI:	inbrowser.ai / @InBrowserAI

Additional slides

Open-source community

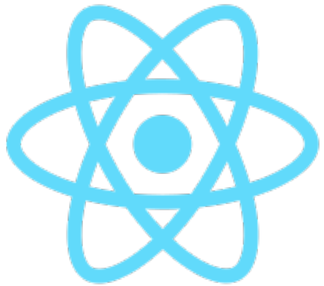

InBrowser.AI

https://inbrowser.ai/
https://twitter.com/InBrowserAI

Pre-training

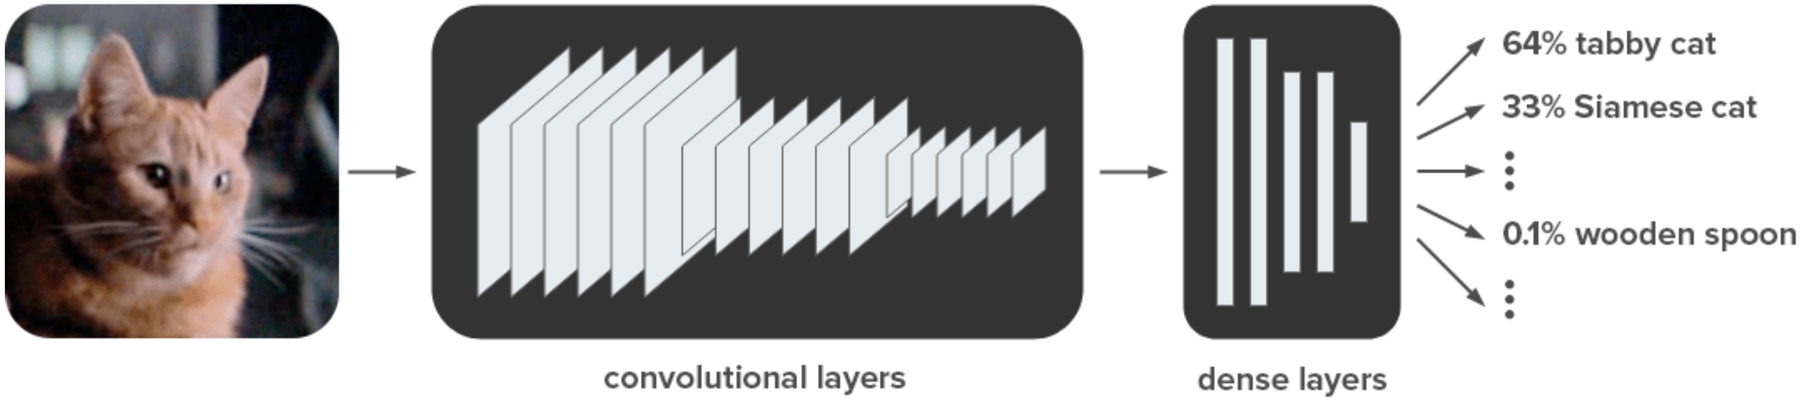

Transfer learning

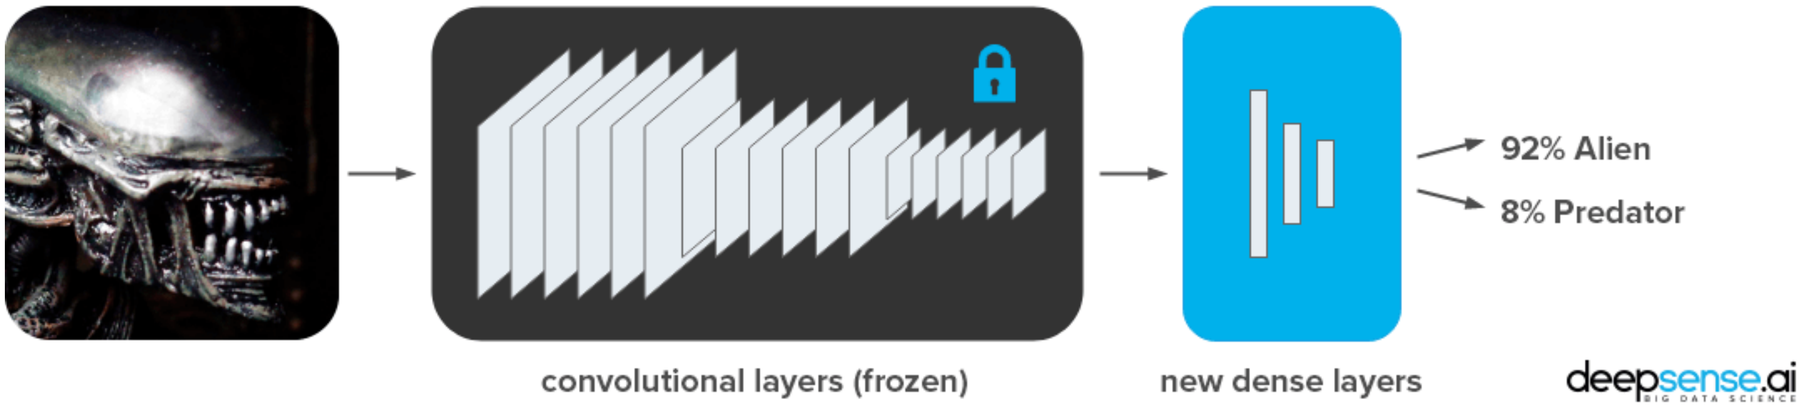

https://deepsense.ai/keras-vs-pytorch-avp-transfer-learning/

Machine Learning events

PyData

PyData Warsaw 2018

19-20 November 2018

https://pydata.org/warsaw2018/

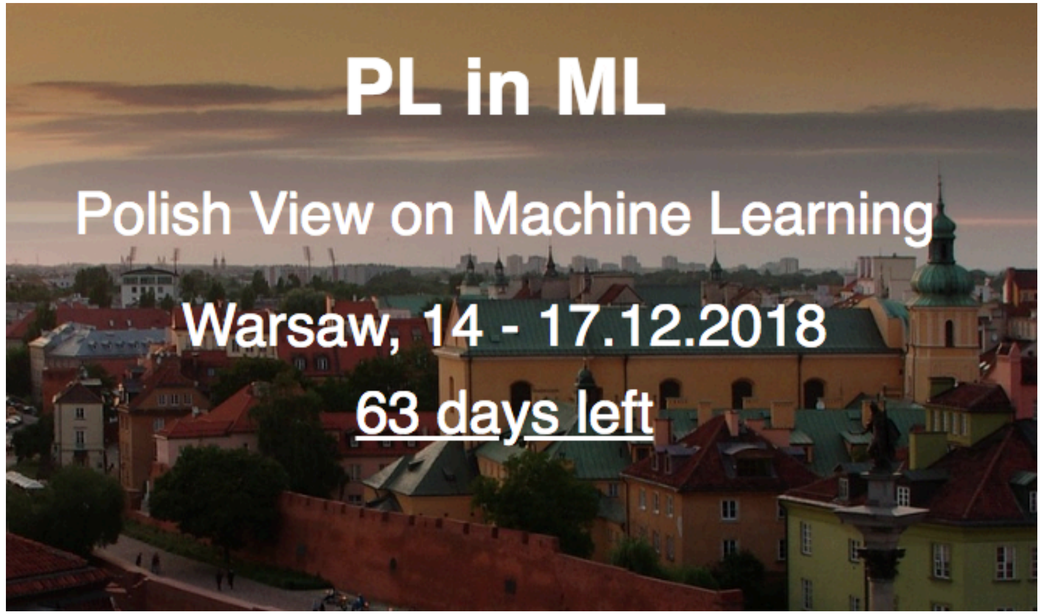

http://plinml.mimuw.edu.pl/

Services accessible via API

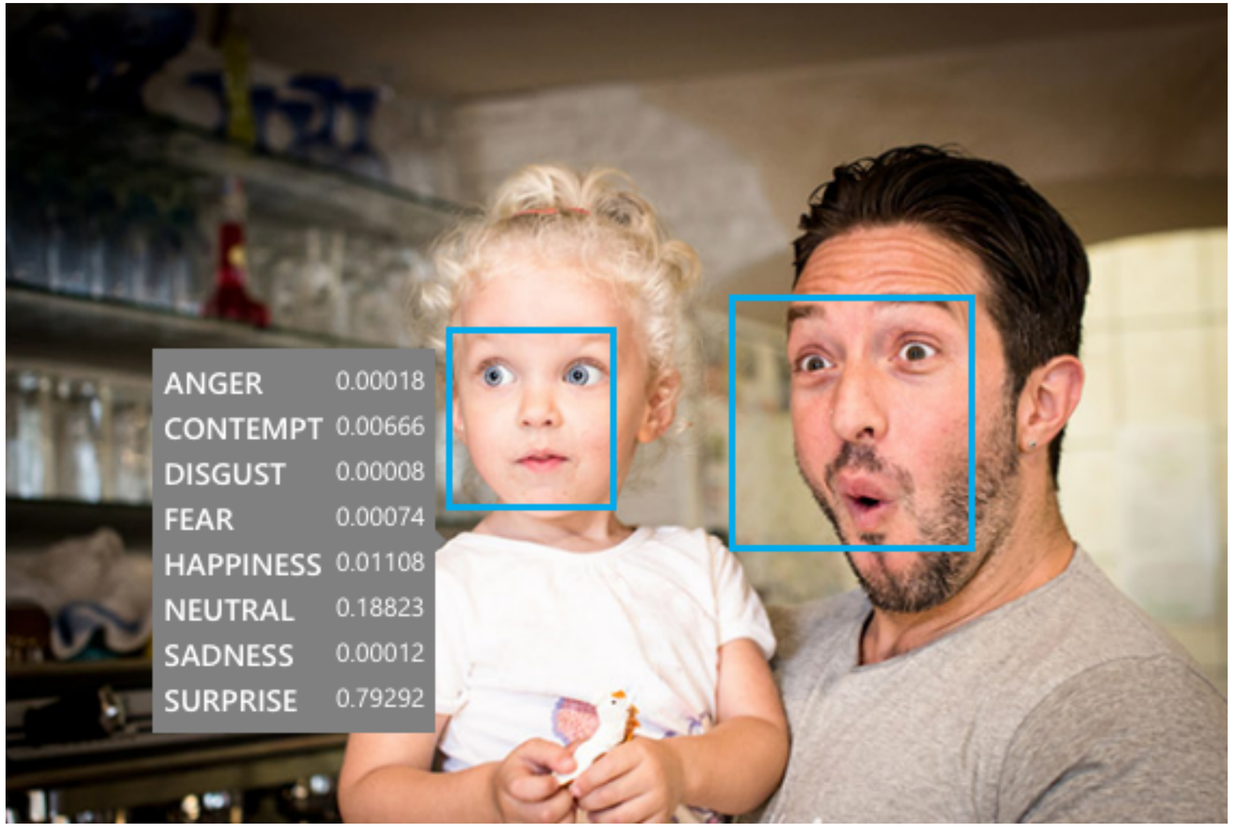


```
Detection result:
2 faces detected

 JSON:
 [
  {
    "faceRectangle": {
      "top": 141,
      "left": 356,
      "width": 123,
      "height": 123
    },
    "scores": {
      "anger": 2.45541414E-05,
      "contempt": 0.000249862962,
      "disgust": 4.51454434E-05,
      "fear": 0.001153976,
      "happiness": 5.58101638E-05,
      "neutral": 0.004939176,
      "sadness": 1.091316E-05,
      "surprise": 0.993520558
    }
  },
```

https://azure.microsoft.com/en-us/services/cognitive-services/face/#recognition
https://goberoi.com/comparing-the-top-five-computer-vision-apis-98e3e3d7c647

Machine learning models

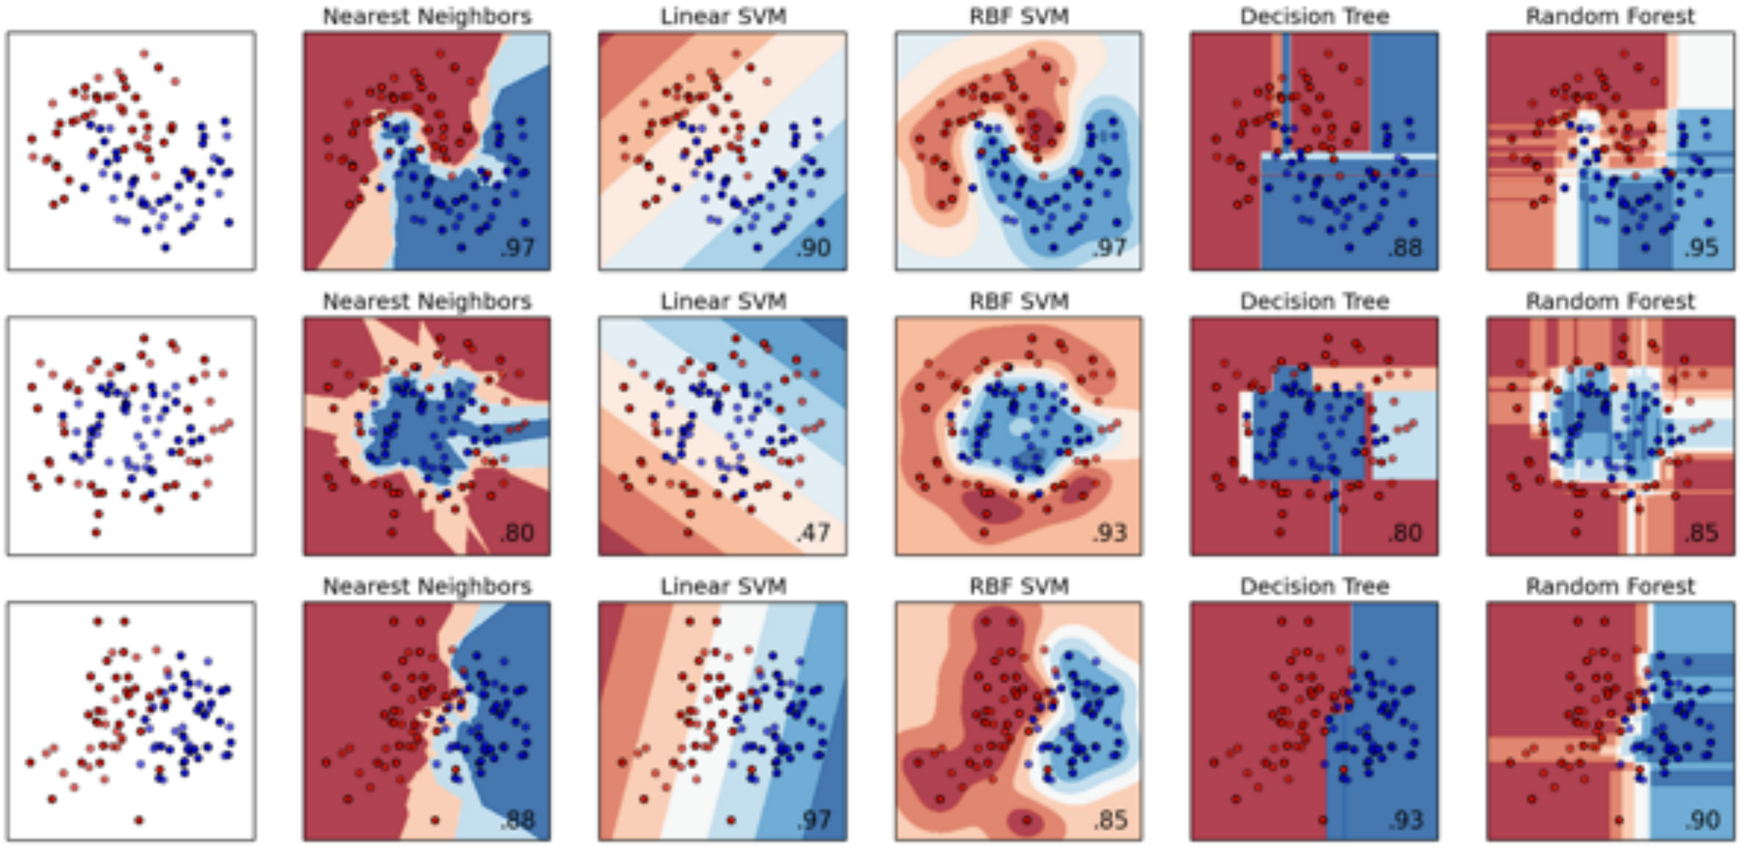

scikit-learn.org/stable/auto_examples/classification/plot_classifier_comparison.html

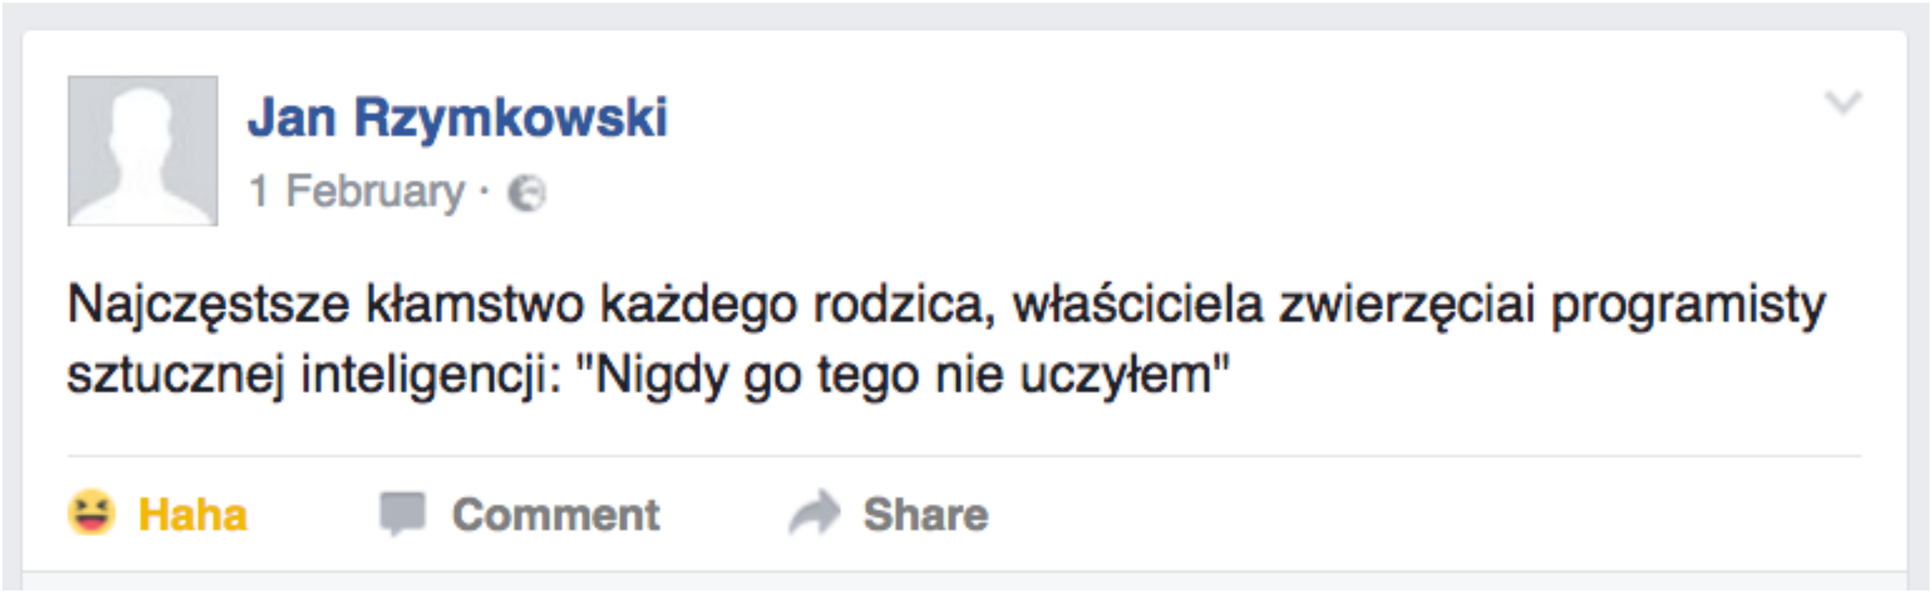

The biggest lie of every
parent, pet owner
and neural network trainer:

“I’ve never shown it that,
it must have learned it by itself!”

Science-based games

https://github.com/stared/science-based-games-list

Thank you!
Questions?

p.migdal.pl - blog
pmigdal@gmail.com @pmigdal github/stared

- inbrowser.ai / Twitter: @InBrowserAI
- Simple diagrams of convolved neural networks
- Keras or PyTorch as your first deep learning framework
- Learning Deep Learning with Keras
- king - man + women is queen; but why?

InBrowser.AI

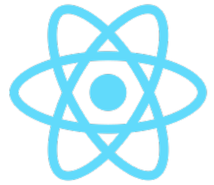

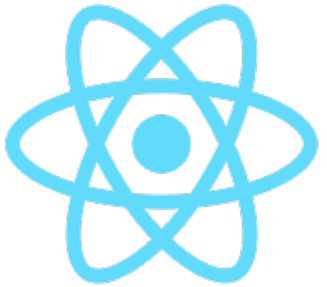

InBrowser.AI

https://inbrowser.ai/
https://twitter.com/InBrowserAI